

NetBackup™ Web UI Administrator's Guide

Release 11.0

NetBackup™ Web UI Administrator's Guide

Last updated: 2025-03-25

Legal Notice

Copyright © 2025 Cohesity, Inc. All rights reserved.

Cohesity, Veritas, the Cohesity Logo, Veritas Logo, Veritas Alta, Cohesity Alta, and NetBackup are trademarks or registered trademarks of Cohesity, Inc. or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Cohesity is required to provide attribution to the third party ("Third-party Programs"). Some of the Third-party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the Third-party Legal Notices document accompanying this Cohesity product or available at:

<https://www.veritas.com/about/legal/license-agreements>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Cohesity, Inc. and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. Cohesity, Inc. SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Cohesity as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Cohesity, Inc.
2625 Augustine Drive
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Cohesity account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Cohesity website:

<https://sort.veritas.com/documents>

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

NB.docs@veritas.com

You can also see documentation information or ask a question on the Cohesity community site:

<http://www.veritas.com/community/>

Cohesity Services and Operations Readiness Tools (SORT)

Cohesity Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Section 1	About NetBackup	24
Chapter 1	Introducing NetBackup	25
	About NetBackup	25
	NetBackup web UI features	27
	NetBackup documentation	29
	NetBackup administration interfaces	29
	About security certificates for NetBackup hosts	30
	First-time sign in to the NetBackup web UI	31
	Sign in to the NetBackup web UI	32
	Sign out of the NetBackup web UI	34
	Using the NetBackup web UI	35
	Terminology	38
Chapter 2	Administering NetBackup licenses	41
	About NetBackup licenses	41
	Add licenses	42
	View licenses	43
	Renew licenses	43
	Remove licenses	43
Section 2	Monitoring and notifications	45
Chapter 3	Monitoring NetBackup activity	46
	The NetBackup dashboard	46
	Activity monitor	48
	Monitor NetBackup daemons	48
	Monitor NetBackup processes	49
	Job monitoring	49
	Workloads that require a custom RBAC role for specific job permissions	50
	View a job	52
	View the jobs in the List view	52

	View the jobs in the Hierarchy view	53
	Jobs: cancel, suspend, restart, resume, delete	53
	View the logs for a job	54
	Search for or filter jobs in the jobs list	54
	Create a jobs filter	55
	Edit, copy, or delete a jobs filter	57
	Import or export job filters	59
	View the status of a redirected restore	59
	Troubleshooting the viewing and managing of jobs	60
Chapter 4	Device monitor	63
	About the Device Monitor	63
	About media mount errors	64
	About pending requests and actions	65
	About pending requests for storage units	66
	Resolve a pending request	66
	Resolve a pending action	67
	Resubmit a pending request	68
	Deny a pending request	68
Chapter 5	Notifications	69
	Job notifications	69
	Send email notifications for job failures	69
	Send notifications to the backup administrator about failed backups	72
	Send notifications to a host administrator about backups	73
	Configure the nbmail.cmd script on the Windows hosts	73
	NetBackup event notifications	75
	View notifications	76
	Modify or disable NetBackup event notifications in the web UI	76
	NetBackup event types supported with notifications	78
	About configuring automatic notification cleanup tasks	83
Chapter 6	Registering the data collector	85
	About the data collector	85
	Register the data collector with Cohesity Alta View	86
	Renew Cohesity Alta View token	86
	Register the data collector with NetBackup IT Analytics	87
	View and modify the data collector registration	88
	Unregister the data collector	89

Section 3	Configuring hosts	90
Chapter 7	Managing host properties	91
	Overview of host properties	93
	View or edit the host properties of a server or client	94
	Host information and settings in Host properties	95
	Reset a host's attributes	96
	Active Directory properties	97
	Backup pool host properties	97
	Busy file settings properties	99
	Activating the Busy file settings in host properties	100
	Clean up properties	101
	Client name properties	103
	Client attributes properties	104
	General tab of the Client attributes properties	106
	Connect options tab of the Client attributes properties	109
	Windows open file backup tab of the Client attributes properties	110
	Client settings properties for UNIX clients	113
	VxFS file change log (FCL) for incremental backups property	115
	Client settings properties for Windows clients	117
	How to determine if change journal support is useful in your	
	NetBackup environment	120
	Guidelines for enabling NetBackup change journal support	120
	Cloud Storage properties	121
	Credential access properties	122
	Data Classification properties	122
	Add a data classification	123
	Default job priorities properties	124
	Understanding the job priority setting	125
	Distributed application restore mapping properties	126
	Encryption properties	127
	Additional encryption methods for Windows clients	128
	Enterprise Vault properties	129
	Enterprise Vault hosts properties	130
	Exchange properties	131
	About the Exchange credentials in the client host properties	132
	Exclude list properties	133
	Add an entry to an exclude list	134
	Add an exception to the exclude list	134
	Syntax rules for exclude lists	135

About creating an include list on a UNIX client	137
Traversing excluded directories	138
Fibre transport properties	139
About Linux concurrent FT connections	141
Firewall properties	142
General server properties	144
Forcing restores to use a specific server	146
Global attributes properties	147
About constraints on the number of concurrent jobs	150
Setting up mailx email client	150
Logging properties	151
Logging levels	153
Lotus Notes properties	155
Media properties	157
Results when media overwrites are not permitted	160
Recommended use for Enable SCSI reserve property	161
Network properties	162
Network settings properties	162
Reverse host name lookup property	163
Use the IP address family property	164
Nutanix AHV access hosts	164
Port ranges properties	165
Registered ports and dynamically-allocated ports	166
Preferred network properties	166
Add or edit a Preferred network setting	169
How NetBackup uses the directives to determine which network to use	170
Configurations to use IPv6 networks	173
Configurations to use IPv4 networks	175
Order of directive processing in the Preferred network properties	176
bptestnetconn utility to display Preferred network information	177
Configuration to prohibit using a specified address	178
Configuration to prefer a specified address	179
Configuration that restricts NetBackup to one set of addresses	180
Configuration that limits the addresses, but allows any interfaces	181
Properties setting in host properties	181
RHV access hosts properties	182
Resilient network properties	182
View the resiliency status of a client	184

About Resilient jobs	185
Resilient connection resource usage	185
Specify resilient connections for clients	186
Resource limit properties	187
Restore failover properties	188
Assigning an alternate media server as a failover restore server	189
Retention periods properties	189
Changing a retention period	191
Determining retention periods for volumes	192
Retention Periods with end dates beyond 2038, excluding Infinity	192
Scalable Storage properties	193
Configuring advanced bandwidth throttling settings	194
Advanced bandwidth throttling settings	195
Servers properties	197
Add a server to a servers list	198
Remove a server from a servers list	199
Enable inter-node authentication for a NetBackup clustered primary server	199
Changing the primary server that performs backups and restores for a client	200
SharePoint properties	201
Consistency check options for SharePoint Server	202
SLP settings properties	202
About batch creation logic in Storage Lifecycle Manager	207
Throttle bandwidth properties	208
Timeouts properties	209
Universal settings properties	212
UNIX client properties	214
UNIX Server properties	214
User account settings properties	214
VMware access hosts properties	215
Windows client properties	216
Configuration options not found in the host properties	216
About using commands to change the configuration options on UNIX or Linux clients and servers	216

Chapter 8

Managing credentials for workloads and systems that NetBackup accesses	218
Overview of credential management in NetBackup	218
Adding credentials in NetBackup	219

	Add a credential for NetBackup Callhome Proxy	219
	Add a credential for an external KMS	220
	Add a credential for Network Data Management Protocol (NDMP)	221
	Edit or delete a named credential	222
	Edit or delete Network Data Management Protocol (NDMP) credentials in NetBackup	223
	Add a configuration for an external CMS server	223
	Configure external credentials	224
	Add a credential for CyberArk	225
	Edit or delete the configuration for an external CMS server	227
	Troubleshooting the external CMS server issue	228
Chapter 9	Managing deployment	229
	About the deployment policies utility	229
	Managing the NetBackup Package repository	230
	Update host	231
	Deployment policies	232
	Attributes tab in Deployment management	232
	Hosts tab in Deployment management	233
	Schedules tab in Deployment management	234
	Security options tab in Deployment management	235
	Copy a deployment policy	237
	Manually deploy a deployment policy	237
	Deployment job status	238
Section 4	Configuring storage	240
Chapter 10	Overview of storage options	241
	About storage configuration	241
Chapter 11	Configuring disk storage	243
	Create a Media Server Deduplication Pool storage server	243
	Integrating MSDP Cloud and CMS	245
	Migrating or updating MSDP Cloud and CMS	248
	Create a Media Server Deduplication Pool (MSDP) storage server for image sharing	248
	Create an AdvancedDisk, OpenStorage (OST), or Cloud Connector storage server	250
	Create an MSDP server for MSDP volume group (MVG)	252
	Create the MVG volume	253

Edit a storage server	254
About configuring disk pool storage	254
Create a disk pool	255
Edit a disk pool	257
Share images from an on-premises location to the cloud	257
Overview of universal shares	258
About the MSDP object store	258
Configuring the MSDP object store	259
Resetting the MSDP object store root user credentials	259
 Chapter 12 Managing media servers	261
Add a media server	261
Activate or deactivate a media server	262
Stop or restart the media device manager	263
About NetBackup server groups	263
Add a server group	264
Delete a server group	264
 Chapter 13 Configuring storage units	266
Overview of storage units	266
About configuring BasicDisk storage	267
Create a storage unit	268
Edit storage unit settings	269
Copy a storage unit	270
Delete a storage unit	271
Tape storage unit considerations	272
Disk storage unit considerations	273
About the disk storage model	274
Configure the NetBackup service credentials for CIFS storage and disk storage units	275
Disk storage units in storage lifecycle policies	275
Maintain the available disk space on disk storage units	276
NDMP storage unit considerations	277
 Chapter 14 Configuring robots and tape drives	278
NetBackup robot types	278
Prerequisite for configuring robots and drives	279
About configuring robots and tapes drives in NetBackup	280
About drive name rules	280
Configure drives and robots by using the wizard	281
Configure drive name rules	283

Updating the device configuration by using the wizard	284
Robot properties and configuration options	284
Robot control (robot configuration options)	285
Managing robots	287
Change the robot control properties of a robot	287
Delete a robot	288
Managing tape drives	288
Change a drive comment	289
About downed drives	289
Change a drive operating mode	290
Change the operating mode for a drive path	290
Clean a tape drive	291
Delete a drive	291
Reset a drive	292
Reset the mount time of a drive	292
Set the drive cleaning frequency	293
View drive details	293

Chapter 15 Configuring tape media 294

About NetBackup tape volumes	294
About NetBackup volume pools	295
About NetBackup volume groups	296
NetBackup media types	297
About adding volumes	298
About adding robotic volumes	299
About adding standalone volumes	299
Add a volume	300
Volume properties	300
Managing volumes	304
Edit a volume	304
About moving volumes	304
Move volumes	305
About recycling a volume	306
About assigning and deassigning volumes	307
Delete a volume	308
Changing the media owner of a volume	309
Changing the volume group assignment	309
About rules for moving volumes between groups	310
Rescan and update barcodes	310
About barcode rules	311
About injecting and ejecting volumes	311
Label a volume	313

Erase a volume	314
Freeze or unfreeze a volume	315
Suspend or unsuspend volumes	316
Managing volume pools	316
Add a volume pool	316
Edit or delete a volume pool	317
Volume pool properties	318
Managing volume groups	319
Delete a volume group	319
Move a volume group	319

Chapter 16 Inventorying robots 321

About robot inventory	322
When to inventory a robot	322
About showing a robot's contents	324
About inventory results for API robots	325
Show the media in a robot	325
About comparing a robot's contents with the volume configuration	326
Comparing media in a robot with the volume configuration	326
About previewing volume configuration changes	327
Previewing volume configuration changes for a robot	327
About updating the NetBackup volume configuration	328
Update the NetBackup volume configuration with a robot's contents	329
Robot inventory options	329
Advanced options for robot inventory settings	331
Configure media ID generation rules	334
Barcode rules settings	335
Media ID generation options	337
Configure media settings	338
About media type mapping rules	339
Configure media type mappings	340

Chapter 17 Staging backups 341

About staging backups	341
About basic disk staging	342
Create a BasicDisk storage unit with disk staging	343
Disk staging storage unit size and capacity	344
Finding potential free space on a BasicDisk disk staging storage unit	345
Schedule settings for disk staging	347

Chapter 18	Troubleshooting storage configuration	351
	Registering a media server	351
	Storage configuration issues	352
Section 5	Configuring backups	353
Chapter 19	Overview of backups in the NetBackup web UI	354
	Backup methods supported in the NetBackup web UI	354
	Policy vs. protection plan FAQs	355
	Support for NetBackup classic policies	355
	Supported protection plan types	356
Chapter 20	Managing classic policies	357
	Add a policy	357
	Example policy - Exchange Server DAG backup	358
	Example policy - Sharded MongoDB cluster	359
	Example policy - Epic-Large-File	361
	Edit, copy, or delete a policy	362
	Deactivate or activate a policy	363
	View automanaged policies and SLPs	364
	About automanaged policies or storage lifecycle policies	364
	Perform manual backups	365
	About the Epic-Large-File policy type	366
Chapter 21	Managing protection plans	368
	Create a protection plan	368
	Customizing protection plans	374
	Edit or delete a protection plan	375
	Subscribe an asset or an asset group to a protection plan	376
	Unsubscribe an asset from a protection plan	377
	View protection plan overrides	377
	Copy a protection plan policy (automanaged policy) to a classic policy	378
	About Backup now	379
Chapter 22	Protecting the NetBackup catalog	382
	About the NetBackup catalog	382
	Catalog backups	383

The catalog backup process	383
Prerequisites for backing up the NetBackup catalog	384
Configuring catalog backups	385
Backing up NetBackup catalogs manually	386
Concurrently running catalog backups with other backups	387
Catalog policy schedule considerations	387
How catalog incrementals and standard backups interact on UNIX	388
Determining whether or not a catalog backup succeeded	388
Strategies that ensure successful NetBackup catalog backups	389
Disaster recovery emails and the disaster recovery files	389
Disaster recovery packages	390
Set the passphrase to encrypt disaster recovery packages	391
Recovering the catalog	393

Chapter 23 Managing backup images 394

About the Catalog utility	394
Catalog utility search criteria and backup image details	395
Verify backup images	397
Promote a copy to a primary copy	398
Duplicate backup images	399
Multiplexed duplication considerations	403
Jobs that appear while making multiple copies	403
Expire backup images	403
About importing backup images	404
About importing expired images	404
Import backup images, Phase I	405
Import backup images, Phase II	406

Chapter 24 Pausing data protection activity 408

Pause backups and other activity	408
Allow the automatic pause of data protection activity	409
Pause backups and other activity on a client	409
View paused backups and other paused activities	409
Resume data protection activity	410

Section 6	Managing security	411
Chapter 25	Security events and audit logs	412
	View security events and audit logs	412
	About NetBackup auditing	413
	User identity in the audit report	416
	Audit retention period and catalog backups of audit records	417
	Viewing the detailed NetBackup audit report	417
	Send audit events to system logs	420
	Send audit events to log forwarding endpoints	421
Chapter 26	Managing security certificates	422
	About security management and certificates in NetBackup	422
	NetBackup host IDs and host ID-based certificates	423
	Manage NetBackup security certificates	424
	Reissue a NetBackup certificate	425
	Manage NetBackup certificate authorization tokens	427
	Using external security certificates with NetBackup	428
	Configure an external certificate for the NetBackup web server	
		429
	Remove the external certificate configured for the web server	
		430
	Update or renew the external certificate for the web server	431
	View external certificate information for the NetBackup hosts in	
	the domain	431
Chapter 27	Managing host mappings	433
	View host security and mapping information	433
	Approve or add mappings for a host that has multiple host names	
		434
	Example host mappings	436
	Remove mappings for a host that has multiple host names	439
Chapter 28	Minimizing security configuration risk	441
	About security configuration risk	441
	Security settings to be configured to minimize risk	443
	Set the current posture as security baseline	445
	Manage security baseline	445
	Manage security baseline from Alta View UI	446

Chapter 29	Configuring multi-person authorization	447
	About multi-person authorization	447
	Workflow to configure multi-person authorization for NetBackup operations	450
	RBAC roles and permissions for multi-person authorization	451
	Multi-person authorization process with respect to roles	452
	NetBackup operations that need multi-person authorization	455
	Configure multi-person authorization	457
	View multi-person authorization tickets	458
	Manage multi-person authorization tickets	459
	Add exempted users	459
	Schedule expiration and purging of multi-person authorization tickets	460
	Disable multi-person authorization	460
Chapter 30	Managing user sessions	462
	Terminate a NetBackup user session	462
	Unlock a NetBackup user	463
	Configure when idle sessions should time out	464
	Configure the maximum of concurrent user sessions	464
	Configure the maximum of failed sign-in attempts	464
	Display a banner to users when they sign in	465
Chapter 31	Configuring multifactor authentication	467
	About multifactor authentication	467
	Configure multifactor authentication for your user account	468
	Disable multifactor authentication for your user account	469
	Enforce multifactor authentication for all users	469
	Configure multifactor authentication for your user account when it is enforced in the domain	470
	Reset multifactor authentication for a user	470
Chapter 32	Managing the global security settings for the primary server	472
	View the Certificate authority for secure communication	473
	Disable communication with NetBackup 8.0 and earlier hosts	473
	Disable automatic mapping of NetBackup host names	474
	Configure the global data-in-transit encryption setting	474
	About NetBackup certificate deployment security levels	475
	Select a security level for NetBackup certificate deployment	478

About TLS session resumption	478
Set a passphrase for disaster recovery	479
Validate the disaster recovery package passphrase	480
About trusted primary servers	480
About the certificate to use to add a trusted primary server	481
Add a trusted primary server	482
Remove a trusted primary server	483
Configure the audit retention period	484

Chapter 33	Using access keys, API keys, and access codes	485
	Access keys	485
	API keys	485
	Add an API key or view API key details (Administrators)	486
	Edit, reissue, or delete an API key (Administrators)	487
	Add an API key or view your API key details	489
	Edit, reissue, or delete your API key	489
	Use an API key with NetBackup REST APIs	490
	Access codes	491
	Request CLI access through web UI authentication	491
	Approve the CLI access request of another user	492
	Edit the settings for command-line access	493

Chapter 34	Configuring authentication options	494
	Sign-in options for the NetBackup web UI	494
	Configure user authentication with smart cards or digital certificates	
		495
	Configure smart card authentication with a domain	495
	Configure smart card authentication without a domain	496
	Edit the configuration for smart card authentication	497
	Add or delete a CA certificate that is used for smart card authentication	498
	Disable or temporarily disable smart card authentication	499
	About single sign-on (SSO) configuration	499
	Configure NetBackup for single sign-on (SSO)	501
	Configure the SAML KeyStore	502
	Configure the SAML keystore and add and enable the IDP configuration	505
	Enroll the NetBackup primary server with the IDP	508
	Manage an IDP configuration	509
	Video: Configure single sign-on in NetBackup	511
	Troubleshooting SSO	512

	Redirection issues	512
	Unable to sign in due to authorization-related issues	513
Chapter 35	Managing role-based access control	516
	RBAC features	516
	Authorized users	517
	Configuring RBAC	517
	Notes for using NetBackup RBAC	518
	Add AD or LDAP domains	519
	View users in RBAC	519
	Add a user to a role (non-SAML)	519
	Add a smart card user to a role (non-SAML, without AD/LDAP)	520
	Add a user to a role (SAML)	521
	Remove a user from a role	522
	Default RBAC roles	522
	Add a custom RBAC role	525
	Edit or remove a role a custom role	526
	Add a custom RBAC role to restore Azure-managed instances	527
	Add a custom RBAC role for a PaaS administrator	529
	Add a custom RBAC role for a Malware administrator	530
	Role permissions	530
	Manage access permission	531
	View access definitions	533
Chapter 36	Disabling access to NetBackup interfaces for OS Administrators	534
	Disable command-line (CLI) access for operating system (OS) administrators	534
	Disable web UI access for operating system (OS) administrators	535
Section 7	Detection and reporting	536
Chapter 37	Detecting anomalies	537
	About backup anomaly detection	537
	How a backup anomaly is detected	538
	Configure backup anomaly detection settings	539
	View backup anomalies	542
	Disable backup anomaly detection and computation of entropy and file attributes for a client	543

	About system anomaly detection	544
	Configure system anomaly detection settings	544
	Configure rules-based anomaly detection	545
	Configure risk engine-based anomaly detection	546
	View system anomalies	550
Chapter 38	Malware scanning	551
	About malware scanning	551
	Workflow for malware scanning	554
	Configuring a scan host pool	560
	Prerequisites for scan host pool	561
	Configure a new scan host pool	561
	Add a new host in a scan host pool	562
	Managing a scan host	563
	Add an existing scan host	563
	Validating the scan host pool configuration	564
	Remove the scan host	564
	Activate/Deactivate the scan host	565
	Managing credentials for malware scanning	565
	Configure resource limits for malware detection	567
	Perform a malware scan	568
	Scanning backup images	569
	Assets by policy type	572
	Assets by workload type	573
	Managing scan tasks	574
	View the malware scan status	574
	Actions for malware scanned images	576
	Recover from malware-affected images (clients protected by policies)	579
	Recover from malware-affected images (clients protected by protection plan)	580
	Clean file recovery for virtual workload (VMware)	581
Chapter 39	Usage reporting and capacity licensing	584
	Track protected data size on your primary servers	584
	Add a local primary server	585
	View license types in usage reporting	586
	Download usage reports	586
	Scheduling reports for capacity licensing	587
	Other configuration for incremental reporting	591
	Troubleshooting failures for usage reporting and incremental reporting	593

Chapter 40	Reports	595
	About the reports utility	595
	Run a report	595
	Copy a report text to another document	596
Section 8	NetBackup workloads and NetBackup Flex Scale	597
Chapter 41	NetBackup SaaS Protection	598
	Overview of NetBackup for SaaS	598
	Adding NetBackup SaaS Protection Hubs	600
	Configuring the autodiscovery frequency	601
	Viewing asset details	601
	Configuring permissions	602
	Troubleshooting SaaS workload issues	603
Chapter 42	NetBackup Flex Scale	605
	Managing NetBackup Flex Scale	605
	Access NetBackup from the Flex Scale infrastructure management console	606
	Manage NetBackup and the NetBackup Flex Scale cluster management from the NetBackup Flex Scale web UI	607
	Access NetBackup Flex Scale from the NetBackup web UI	608
Chapter 43	NetBackup workloads	610
	Protection of other asset types and clients	610
Section 9	Administering NetBackup	611
Chapter 44	Management topics	612
	Configuring the NetBackup Client Service	612
	Units of measure used with NetBackup	613
	NetBackup naming conventions	614
	Wildcard use in NetBackup	615
Chapter 45	Managing client backups and restores	618
	About server-directed restores	618
	About client-redirected restores	620

About restore restrictions	621
Allowing all clients to perform redirected restores	621
Allowing a single client to perform redirected restores	622
Allowing redirected restores of a specific client's files	622
Examples of redirected restores	623
About restoring the files that have Access Control Lists (ACLs)	629
About setting the original atime for files during restores on UNIX	630
Restoring the System State	630
About the backup and restore of compressed files on VxFS file systems	632
About backups and restores on ReFS	633

Section 10 Disaster recovery and troubleshooting

.....	634
Chapter 46 Disaster recovery of NetBackup	635
About disaster recovery of NetBackup	635
Chapter 47 Managing Resiliency Platforms	636
About Resiliency Platform in NetBackup	636
Understanding the terms	637
Configuring a Resiliency Platform	638
Add a Resiliency Platform	638
Configure a third-party CA certificate	639
Edit or delete a Resiliency Platform	639
View the automated or not-automated VMs	640
Troubleshooting NetBackup and Resiliency Platform issues	642
Chapter 48 Managing Bare Metal Restore (BMR)	644
About Bare Metal Restore (BMR)	644
Add a custom role for a Bare Metal Restore (BMR) administrator	645
Chapter 49 Troubleshooting the NetBackup Web UI	647
Tips for accessing the NetBackup web UI	647
If a user doesn't have the correct permissions or access in the NetBackup web UI	649
Unable to validate the user or group when configuring LDAP server	649

Section 11	Other topics	650
Chapter 50	Additional NetBackup catalog information	651
	Parts of the NetBackup catalog	651
	NetBackup databases and configuration files	652
	About the NetBackup image database	654
	About the catalog backup of cloud configuration files	656
	Archiving the catalog and restoring from the catalog archive	657
	Enabling intelligent catalog archiving (ICA) to reduce the number of .f files	660
	Creating a catalog archiving policy	664
	Catalog archiving commands	665
	Catalog archiving considerations	667
	Extracting images from the catalog archives	668
	Estimating catalog space requirements	668
	NetBackup file size considerations on UNIX systems	670
	Moving the image catalog	670
	About image catalog compression	672
	About the file hash search in NetBackup	675
	Configuring the file hash server	676
	Enabling the file hash server on the NetBackup primary server	677
	Calculating the file hash	677
	Searching the files using the file hash	678
	Identifying the backups that have the file hash enabled	679
	Removing the file hash from the backup	679
	Migrating the file hash data from one server to another	679
	Configuring the backup of file hash data on the file hash server	680
	Restoring the file hash data to the file hash server	681
Chapter 51	About the NetBackup database	682
	About the NetBackup database installation	682
	About NetBackup primary server installed directories and files	682
	NetBackup configuration entry	685
	NetBackup database server management	686
	The NetBackup database and clustered environments	687
	Post-installation tasks	687
	Changing the NetBackup database password	688
	Moving a database after installation	689

Copying the NetBackup databases	691
Creating the NBDB database manually	691
Using the NetBackup Database Administration utility on Windows	693
General tab of the NetBackup Database Administration utility	694
Tools tab of the NetBackup Database Administration utility	695
Using the NetBackup Database Administration utility on UNIX	698
Select/Restart Database and Change Password menu options	700
Database Space Management menu options	700
Database Validation Check and Rebuild menu options	701
Move Database menu options	702
Unload Database menu options	703
Backup and Restore Database menu options	703

About NetBackup

- [Chapter 1. Introducing NetBackup](#)
- [Chapter 2. Administering NetBackup licenses](#)

Introducing NetBackup

This chapter includes the following topics:

- [About NetBackup](#)
- [NetBackup web UI features](#)
- [NetBackup documentation](#)
- [NetBackup administration interfaces](#)
- [Using the NetBackup web UI](#)
- [Terminology](#)

About NetBackup

NetBackup provides a complete, flexible data protection solution for a variety of platforms. The platforms include Windows, UNIX, and Linux systems.

NetBackup administrators can set up periodic or calendar-based schedules to perform automatic, unattended backups for clients across a network. An administrator can carefully schedule backups to achieve systematic and complete backups over a period of time, and optimize network traffic during off-peak hours. The backups can be full or incremental: Full backups back up all indicated client files, while incremental backups back up only the files that have changed since the last backup.

The NetBackup administrator can allow users to back up, restore, or archive the files from their computer. (An archive operation backs up a file, then deletes it from the local disk if the backup is successful.)


NetBackup includes both the server and the client software as follows:

- Server software resides on the computer that manages the storage devices.

- Client software resides on computers that contain data to back up. (Servers also contain client software and can be backed up.)

Figure 1-1 shows an example of a NetBackup storage domain.

Figure 1-1 NetBackup storage domain example



NetBackup accommodates multiple servers that work together under the administrative control of one NetBackup primary server in the following ways:

- The primary server manages backups, archives, and restores. The primary server is responsible for media and device selection for NetBackup. Typically, the primary server contains the NetBackup catalog. The catalog contains the internal databases that contain information about NetBackup backups and configuration.
- Media servers provide additional storage by allowing NetBackup to use the storage devices that are attached to them. Media servers can also increase performance by distributing the network load. Media servers can also be referred to by using the following terms:
 - Device hosts (when tape devices are present)
 - Storage servers (when I/O is directly to disk)

- Data movers (when data is sent to independent, external disk devices like OpenStorage appliances)

During a backup or archive, the client sends backup data across the network to a NetBackup server. The NetBackup server manages the type of storage that is specified in the backup policy.

During a restore, users can browse, then select the files and directories to recover. NetBackup finds the selected files and directories and restores them to the disk on the client.

NetBackup web UI features

The NetBackup web user interface provides the following features:

- Ability to access the primary server from a web browser, including Chrome and Firefox. For details on supported browsers for the web UI, see the [NetBackup Software Compatibility List](#).
Note that the NetBackup web UI may behave differently for different browsers. Some functionality, for example a date picker, may not be available on all browsers. These inconsistencies are due to the capabilities of the browser and not because of a limitation with NetBackup.
- A dashboard that displays a quick overview of the information that is important to you.
- Role-based access control (RBAC) that lets the administrator configure user access to NetBackup and to delegate the tasks such as security, storage management, or workload protection.
- Management of NetBackup security settings including certificates, access keys, multi-person authorization, and user sessions.
- Management of hosts including deployment management and NetBackup host properties.
- Management of storage and devices.
- Data protection is achieved through policies or protection plans.
- Detection and reporting features provide for the detection of malware and anomalies and let you track the size of backup data on your primary servers through usage reporting. You can also easily connect to Cohesity NetInsights Console to view and manage NetBackup licensing.

Note: The NetBackup web UI is best viewed at a 1280x1024 or higher screen resolution.

Role-based access control in the NetBackup web UI

NetBackup uses role-based access control to grant access to the web UI. Access control is accomplished through roles.

- A role defines the operations that a user can perform and the access that the user has to any workload assets, protection plans, or credentials. A user can have multiple roles, which allows for full and for flexible customization of user access.
- RBAC is only available for the web UI and the APIs.
Other access control methods for NetBackup are not supported for the web UI and APIs.

Monitor NetBackup jobs and events

The NetBackup web UI lets administrators more easily monitor NetBackup operations and events and identify any issues that need attention.

- The dashboard displays an overview of NetBackup operations and security information. This information includes jobs, certificates, tokens, security events, malware and anomaly detection, and usage reporting.
The dashboard widgets that display depend on a user's RBAC role and permissions.
- Email notifications can be configured so that administrators receive notifications when job failures occur. NetBackup supports any ticketing system that can receive inbound email.

Backup policies

NetBackup classic policies are available for the Administrator that wants to use policies for data protection. Backup policies provide the instructions that NetBackup follows to back up clients, databases, or virtual machines. These instructions include where to store the backup and when and how frequently to perform the backup. For client backups, policies also include the files and directories to back up.

See [“Support for NetBackup classic policies”](#) on page 355.

Protection plans: One place to configure schedules and storage

Data protection with protection plans is fully managed with role-based access control (RBAC). The NetBackup administrator can manage which users can view and manage assets and can perform backups and restores. Each default workload administrator role (for example, Default VMware Administrator) allows a user access to protection plans, jobs, and credentials.

See [“Supported protection plan types”](#) on page 356.

Protection plans offer the following benefits:

- A workload administrator can create and manage protection plans, including the backup schedules and the storage that the plan uses. This administrator selects the protection plans that protect assets.
See [“Role permissions”](#) on page 530.
- In addition to schedules for backups, a protection plan can also include a schedule for replication and long-term retention.
- When you select from your available storage, you can see any additional features available for that storage.
- Users with a workload administrator role can create protection plans, manage credentials, subscribe assets to the protection plans that meet the SLO, and monitor protection status.

Server-directed and self-service recovery

Administrators can perform server-directed restores from the web UI.

See [“About server-directed restores”](#) on page 618.

The workload administrator can perform self-service recovery of VMs, databases, or other asset types. This type of recovery is available for the assets that are protected with recovery points.

For the workloads that support the instant access feature, users can mount a snapshot for immediate access to a VM's files or to a database.

NetBackup documentation

For a complete list of NetBackup technical documents for each supported release, see the *NetBackup Documentation Landing Page* at the following URL:

<https://www.veritas.com/docs/DOC5332>

No responsibility is assumed for the installation and use of the Adobe Acrobat Reader.

NetBackup administration interfaces

NetBackup can be administered with several interfaces. The best choice depends on personal preference and the systems that are available to the administrator.

Table 1-1 NetBackup administration interfaces

Name of interface	Description
NetBackup web user interface	With the NetBackup web user interface (UI), you can view NetBackup activities and manage NetBackup configuration, from a primary server. To start the NetBackup web UI: ■ Users must have a role that is configured for them in NetBackup RBAC. ■ Open a web browser and go to the following URL: <code>https://primaryserver/webui/login</code>
Character-based, menu interface	Run the <code>tpconfig</code> command to start a character-based, menu interface for device management. Use the <code>tpconfig</code> interface from any terminal (or terminal emulation window) that has a <code>termcap</code> or a <code>terminfo</code> definition.
Command line	NetBackup commands are available on both Windows and UNIX platforms. Enter NetBackup commands at the system prompt or use the commands in scripts. All NetBackup administrator programs and commands require root or administrator user privileges by default. For complete information on all NetBackup commands, see the NetBackup Commands Reference Guide.

About security certificates for NetBackup hosts

NetBackup uses security certificates for authentication of NetBackup hosts. The NetBackup security certificates conform to the X.509 Public Key Infrastructure (PKI) standard. A primary server acts as the NetBackup Certificate Authority (CA) and issues NetBackup certificates to hosts.

NetBackup provides two types of NetBackup host security certificates: Host ID-based certificates and host name-based certificates. Host ID-based certificates are based on Universally Unique Identifiers (UUID) that are assigned to each NetBackup host. The NetBackup primary server assigns these identifiers to the hosts.

Any security certificates that were generated before NetBackup 8.0 are now referred to as host name-based certificates. NetBackup is in the process of replacing these older certificates with newer host ID-based certificates. The transition will be completed in future releases and the use of host name-based certificates will be eliminated. However, the transition is ongoing and the current NetBackup version continues to require the older host name-based certificates for certain operations.

NetBackup uses the certificates that are issued from either a NetBackup Certificate Authority or an external certificate authority for host authentication. If you intend to use external certificates on your primary server, you configure the certificates in a post-installation process. The media servers and the clients that use external

certificates can either configure external certificates during the installation or upgrade, or after the installation or upgrade.

First-time sign in to the NetBackup web UI

After the installation of NetBackup, an administrator must sign into the NetBackup web UI from a web browser and create RBAC roles for users. A role gives a user permissions and access to the NetBackup environment through the web UI, based on the user's role in your organization. Some users have access to the web UI by default.

See [“Authorized users”](#) on page 517.

If you do not have access to root or to administrator credentials you can use the `bpbaz -AddRBACPrincipal` command to add an administrator user.

To sign in to a NetBackup primary server using the NetBackup web UI

- 1 Open a web browser and go to the following URL.

`https://primaryserver/webui/login`

The *primaryserver* is the host name or IP address of the NetBackup primary server that you want to sign in to.

If you are not able to access the web UI, refer to [Support and additional configuration](#).

- 2 Enter the administrator credentials and click **Sign in**.

For this type of user	Use this format	Example
Local user	<i>username</i>	jane_doe
Windows user	<i>DOMAIN\username</i>	WINDOWS\jane_doe
UNIX user	<i>username@domain</i>	john_doe@unix

- 3 On the left, select **Security > RBAC**.
- 4 You can give users access to the NetBackup web UI in one of the following ways:
 - Create roles for all users that require access to NetBackup.
 - Delegate the task of creating roles to another user.
Create a role that has permissions to add RBAC roles. This user can then create roles for all users that require access to the NetBackup web UI.

See [“Configuring RBAC”](#) on page 517.

Root or administrator access is no longer needed for the web UI once you have delegated one or more users with permissions to create RBAC roles.

Support and additional configuration

Refer to the following information for help with accessing the web UI.

- Ensure that you are an authorized user.
See [“Authorized users”](#) on page 517.
- For details on supported browsers for the web UI, see the [NetBackup Software Compatibility List](#).
- If port 443 is blocked or in use, you can [configure and use a custom port](#).
- If you want to use an external certificate with the web browser, see the following topic.
See [“Configure an external certificate for the NetBackup web server”](#) on page 429.
- See other tips for accessing the web UI.
See [“Tips for accessing the NetBackup web UI”](#) on page 647.

Sign in to the NetBackup web UI

Authorized users can sign in to a NetBackup primary server from a web browser, using the NetBackup web UI. The NetBackup web user interface (web UI) is available for NetBackup 8.1.2 and later. This interface is available on the primary server and supports the version of NetBackup on that server.

Users should contact their NetBackup security administrator for information on how to sign in.

The following sign-in options are available:

- [Sign in with a username and password](#)
- [Sign in with a certificate or smart card](#)
- [Sign in with single sign-on \(SSO\)](#)

Sign in with a username and password

You can sign in to NetBackup web UI with your username and password.

To sign in to a NetBackup primary server using a username and password

- 1 Open a web browser and go to the following URL.

`https://primaryserver/webui/login`

The *primaryserver* is the host name or IP address of the NetBackup primary server that you want to sign in to.

- 2 Depending on the sign-in options that are available, choose from the following:

- Enter your credentials and click **Sign in**.
- (Conditional) If your user account is configured for multifactor authentication, you are prompted to enter the one-time password. Enter the one-time password and click **Confirm**. See [“About multifactor authentication”](#) on page 467.
- If the default method is not username and password, click **Sign in with username and password**. Then enter your credentials.

The following are example credentials:

For this type of user	Use this format	Example
Local user	<i>username</i>	jane_doe
Windows user	<i>DOMAINusername</i>	WINDOWSjane_doe
UNIX user	<i>username</i>	john_doe

Sign in with a certificate or smart card

You can sign in to NetBackup web UI with a smart card or digital certificate. To use a digital certificate that is not on a smart card, you must first upload the certificate to the browser's certificate manager. See the browser documentation for instructions or contact your certificate administrator for more information.

To sign in with a certificate or smart card

- 1 Open a web browser and go to the following URL.

`https://primaryserver/webui/login`

The *primaryserver* is the host name or IP address of the NetBackup primary server that you want to sign in to.

- 2 Click **Sign in with certificate or smart card**.

If you use the Firefox web browser and have issues with signing in, see this [tech note](#).

- 3 When your browser prompts you, select the certificate.

Sign in with single sign-on (SSO)

You can sign in to NetBackup web UI with the single sign-on (SSO) option if SAML is configured as an identity provider in your NetBackup environment.

To sign in to a NetBackup primary server using SSO

- 1 Open a web browser and go to the following URL.

`https://primaryserver/webui/login`

The *primaryserver* is the host name or IP address of the NetBackup primary server that you want to sign in to.

- 2 Click **Sign in with single sign-on**.

- 3 Follow the steps as provided by your administrator.

On subsequent logons, NetBackup signs you in automatically to the primary server.

Sign out of the NetBackup web UI

Note that NetBackup automatically signs you out of the web UI after 24 hours, which is the maximum time that is allowed for a user session. After that time NetBackup requires that you sign in again. You can also sign out if you want to change the sign-in option that you want to use (username and password, smart card, or single sign-on (SSO)).

To sign out of the NetBackup web UI

- ◆ On the top right, click the profile icon and click **Sign out**.

Using the NetBackup web UI

The **NetBackup web UI** provides an interface for the administrator to manage NetBackup.

Table 1-2 Utilities in the left sidebar of the NetBackup web UI

Item	Description
Dashboard	Displays a quick overview of the information that is important to you.
Activity monitor	Displays NetBackup job information and provides the control over the jobs, services, processes, and drives. See “Activity monitor” on page 48.
Recovery	Administrators can use the Recovery node to perform the following kinds of recovery: ■ Regular recovery - Perform server-directed restores of the assets that are protected by policies. Server-directed restores are currently limited to a subset of policy types. Recovery for a specific workload is performed from the Workloads node. For example, to recover VMware assets go to Workloads > VMware. ■ NetBackup catalog recovery. Recovers a catalog backup in a disaster recovery situation.
Protection	Data protection is achieved through policies or protection plans. See “Backup methods supported in the NetBackup web UI” on page 354.
Workloads	Contains the supported workloads for NetBackup and tools to manage the workload environment, asset credentials, and recovery.
Storage	This node contains the utilities for managing the media and devices that NetBackup uses to store backups.
Catalog	Search for backup images and perform various actions, including: verify the backup contents, duplicate a backup image, promote a copy, expire a backup image, and import a backup image. See “About the Catalog utility” on page 394.

Table 1-2 Utilities in the left sidebar of the NetBackup web UI (*continued*)

Item	Description
Detection and reporting	This node contains the following tools: ■ Anomaly detection - Detects anomalies in backup metadata. See “About backup anomaly detection” on page 537. ■ Malware detection - Finds malware in supported backup images and finds the last good-known image that is malware free. See “About malware scanning” on page 551. ■ Paused protection - Allows NetBackup or authorized users to pause data protection activities. See “Pause backups and other activity” on page 408. ■ Usage - Displays the primary servers that are configured for capacity licensing and their respective consumption details. See “Track protected data size on your primary servers” on page 584. ■ Reporting - Use to compile information to verify, manage, and troubleshoot NetBackup operations. See “About the reports utility” on page 595.
Credential management	Centrally manages the credentials that NetBackup uses to access systems and the workloads that it protects. You can manage credentials for workloads and for systems, client credentials (for NDMP and disk arrays hosts), and External CMS server configurations. See “Overview of credential management in NetBackup” on page 218.
Hosts	Contains the utilities to manage: ■ Deployment management - The main component of VxUpdate that serves as a client or a host upgrade tool. See “Managing the NetBackup Package repository” on page 230. For more information regarding VxUpdate, see the NetBackup Upgrade Guide. ■ Host properties - Use to customize NetBackup configuration options.
Resiliency	Integrates NetBackup and Veritas Resiliency Platform to manage your disaster recovery operations. See “About Resiliency Platform in NetBackup” on page 636.

Table 1-2 Utilities in the left sidebar of the NetBackup web UI (*continued*)

Item	Description
Security	This node contains the utilities to manage settings for security and hosts: ■ Access keys - Provides access the NetBackup interfaces through API keys and access codes. ■ Certificates - Use to manage NetBackup certificates and view external certificates. ■ Host mappings - Use to carry out NetBackup host operations, such as adding or removing host mappings, resetting a host, or generating a reissue token. ■ Multi-person authorization - Ensures that a second authorized user approves actions before they are performed. ■ RBAC - Use predefined or custom RBAC roles to provide NetBackup users with access to NetBackup, based on their role in your organization. ■ Security events - Use to view the sign-in details for NetBackup users and the user-initiated changes that are made to NetBackup. For more information about Security events, see the NetBackup Security and Encryption Guide. ■ Tokens - Manage the authorization tokens in your NetBackup environment. ■ User sessions - Manage the settings for NetBackup user sessions, terminate user sessions, and unlock a user.
Other licensed utilities	Additional licensed utilities appear under the main NetBackup nodes.

In the upper right-hand corner of the web UI are the following settings.

Table 1-3 Utilities in the top toolbar of the NetBackup web UI

Item	Description
Ticket alerts	Displays a summary of any ticket alerts that are available for multiperson authorization.
Notifications	Displays the most recent events that have occurred in the NetBackup environment.
Help	This menu contains links to the NetBackup Help file and the NetBackup APIs.

Table 1-3 Utilities in the top toolbar of the NetBackup web UI (*continued*)

Item	Description
Settings	This menu contains the following settings: ■ Email notifications - Send email notifications when job failures occur. ■ Global security - Configure security settings for the NetBackup domain. ■ Smart card authentication - Map a smart card or certificate for user validation. ■ Data collector registration - Collect metadata from NetBackup to monitor, manage, and report on NetBackup domains. ■ License management - Manage licenses for NetBackup. ■ Guided setup - Guides you through the process to configure storage, discover virtualization and cloud servers, add protection plans, and protect workloads. ■ NetBackup catalog recovery - Recovers a catalog backup in a disaster recovery situation.
Profile	When you click the profile icon, you can see the following information: ■ Current user's sign in attempts. ■ Password expiration date. ■ NetBackup version of the server. ■ The Approve access request option, to approve an access request that you submitted. ■ The Configure multifactor authentication option, to configure multifactor authentication in NetBackup. ■ The options Add API key or View my API key details, to add your own API key or view the details of your existing API key. ■ In the Profile settings you can add an email address to receive notifications and select ticket events for which you want to receive notifications. ■ The Sign out button, to sign out of the web UI.

Terminology

The following table describes the concepts and terms in web user interface.

Table 1-4 Web user interface terminology and concepts

Term	Definition
Administrator	A user that has complete access and permissions to NetBackup and all of the interfaces, including the NetBackup web UI. The root and administrator user have complete access to NetBackup. In the <i>NetBackup Web UI</i> guides, the term <i>NetBackup administrator</i> also refers to a user that has full permissions for NetBackup. Also see <i>role</i> .
Asset group	See <i>intelligent group</i> .
Asset	The data to be protected, such as physical clients, virtual machines, and database applications.
Backup now	An immediate backup of an asset. NetBackup performs a one-time, full backup of an asset using the selected protection plan. This backup does not affect any scheduled backups.
Classic policy	In the NetBackup web UI, indicates that a legacy policy protects the asset.
External certificate	A security certificate that is issued from any CA other than NetBackup.
Intelligent group	Allows NetBackup to automatically select assets for protection based on the criteria (queries) that you specify. An intelligent group automatically stays up-to-date with changes in the production environment. These groups are also referred to as asset groups. These groups appear under the tab Intelligent VM groups or Intelligent groups .
Instant access	Note: Instant access is supported only a select number of workloads and policies. An instant access VM or database that is created from a NetBackup backup image is available almost instantaneously, achieving a near-zero recovery time objective. NetBackup mounts the snapshot directly on the backup storage device and the snapshot is treated as a normal VM or database.
NetBackup certificate	A security certificate that is issued from the NetBackup CA.

Table 1-4 Web user interface terminology and concepts (*continued*)

Term	Definition
Protection plan	A protection plan defines when backups are performed, how long the backups are retained, and the type of storage to use. Once a protection plan is set up, assets can be subscribed to the protection plan.
RBAC	Role-based access control. The role administrator can delegate or limit access to the NetBackup web UI through the roles that are configured in RBAC.
Role	For RBAC, defines the operations that a user can perform and the assets or objects that they can access. For example, you can configure a role to manage recovery of specific databases and the credentials that are needed for backups and restores.
Storage	The storage to which the data is backed up, replicated, or duplicated (for long-term retention).
Subscribe, to a protection plan	The action of selecting an asset or an asset group to subscribe to a protection plan. The asset is then protected according to the schedule in the plan. The web UI also refers to <i>Subscribe</i> as <i>Add protection</i> .
Unsubscribe, from a protection plan	<i>Unsubscribe</i> refers to the action of removing protection or removing an asset or asset group from a plan.
Workload	The type of asset. For example: VMware, Microsoft SQL Server, or Cloud.

Administering NetBackup licenses

This chapter includes the following topics:

- [About NetBackup licenses](#)
- [Add licenses](#)
- [View licenses](#)
- [Renew licenses](#)
- [Remove licenses](#)

About NetBackup licenses

NetBackup uses a common license system that other Cohesity products also use. However, the common license system provides flexibility in the license features that each product implements.

For example, NetBackup does not have a node-locked license system, but some other products do.

Licenses for all purchased NetBackup SKUs must be entered on the primary server. Enter licenses by using one of the following methods:

- During NetBackup primary server installation
The installer prompts you to enter the licenses for all NetBackup products that you plan to install.
You must add either a NetBackup license file or use an in-built evaluation or a temporary production license during primary server installation. For more information, refer to the *NetBackup Installation Guide* or the following article:
https://www.veritas.com/support/en_US/article.100058779

- NetBackup web UI (recommended)
After you install the NetBackup primary server, add the license in the NetBackup web UI. Click **Settings > License management**. Then click **Add license**.
- Command-line interface
After NetBackup primary server installation, use the following command:

```
/usr/opensv/netbackup/bin/admincmd/get_license_key
```


Use the `bpmlicense` command to manage licenses.
Refer to the *NetBackup Commands Reference Guide* for more details.
On UNIX, you can also use the following command:

```
/usr/opensv/netbackup/bin/admincmd/get_license_key
```

Note: Cohesity recommends the use of a browser and the NetBackup web UI to manage licenses remotely.

Add licenses

You can add licenses after primary server installation using the NetBackup web UI.

To add licenses after primary server installation

- 1 In the NetBackup web UI, click **Settings > License management**.
- 2 On the **License management** screen, click **Add license**.
- 3 Add license file using one of the following methods:
 - **Cohesity Entitlement Management System (VEMS)** - Use this method to add the license from the VEMS portal.
 - Sign in to your Cohesity account by specifying the **Username** and **Password**.
 - Select the entitlement that you want to add.
For more information, refer to the [Veritas Entitlement Management System \(VEMS\) User's Guide](#).
 - **File system** - Use this method to add a license file that you have already downloaded on your local host.
 - Click **Browse** to select the `.slf` license file that you want to add.
- 4 Click **Add**.

View licenses

You can view the NetBackup licenses that you have already added, using the web UI.

To view NetBackup licenses

- 1 In the NetBackup web UI, click **Settings > License management**.
- 2 You can see the following license details:
 - **Name** - Name of the license
 - **Status** - Status of the license, such as Active
 - **License type** - Type of license such as Perpetual, Subscription
 - **Activation** - Date when the license was activated
 - **Expiration** - Date when the license will be expired
 - **Entitlement ID** - Unique identification number of each license with respect to the product features it offers and the customer account that is entitled to use it

Renew licenses

You can renew subscription type of licenses.

To renew licenses

- 1 In the NetBackup web UI, click **Settings > License management**.
- 2 Click the **Actions** option for the license that you want to renew.
- 3 Click **Renew**.
- 4 For the VEMS option, enter the username and password.
For the File system option, select the license file.
- 5 Click **Sign in**.
- 6 Click **Renew**.

Remove licenses

You can remove licenses.

To remove licenses

- 1** In the NetBackup web UI, click **Settings > License management**.
- 2** Click the **Actions** option for the license that you want to remove.
- 3** Click **Remove**.

Monitoring and notifications

- [Chapter 3. Monitoring NetBackup activity](#)
- [Chapter 4. Device monitor](#)
- [Chapter 5. Notifications](#)
- [Chapter 6. Registering the data collector](#)

Monitoring NetBackup activity

This chapter includes the following topics:

- [The NetBackup dashboard](#)
- [Activity monitor](#)
- [Job monitoring](#)

The NetBackup dashboard

NetBackup monitors and displays the following information about the NetBackup environment.

Table 3-1 The NetBackup dashboard

Dashboard widget	Description
Jobs	Lists job information, including the number of active and queued jobs and the status of attempted and completed jobs. You can click on a link for specific job details, for example Active jobs. NetBackup opens the Jobs list in the Activity monitor and creates a temporary filter for those jobs on the Jobs tab. ■ If you navigate to another area of the web UI, the filter is removed (if you did not copy and save it). ■ To save the filter, hover over the filter in the toolbar and click Actions > View. Click Copy, make any changes that you want, and then click Save. ■ You can click Actions > Delete to immediately delete the filter.

Table 3-1 The NetBackup dashboard (*continued*)

Dashboard widget	Description
Backup anomaly detection	Displays the total anomalies that are reported so far. See “View backup anomalies” on page 542. Note: An anomalies count of 0 indicates there are no anomalies generated or that the anomaly detection services are not running.
Malware detection	Displays the malware scan result status for the images including Impacted, Not impacted, Failed, In progress, and Pending. <i>(Applicable only when file hash search is configured and domain is not registered with Alta)</i> The View impact analysis report displays the set of backup images impacted, asset client name (to which the backups belong to) and impacted file paths. See “About malware scanning” on page 551.
Paused protection	Lists any paused protection activities for clients. These activities include new backups, duplication, and image expiration. NetBackup pauses protection if it detects malware in backup images. Automatic indicates the activities that are automatically paused by NetBackup. User-initiated indicates an activity that was paused manually by a user. See “Pause backups and other activity” on page 408.
Certificates	Displays the information about the NetBackup host ID-based security certificates and any external certificates in your environment. More details are available in Certificates > External certificates. See “About security management and certificates in NetBackup” on page 422. For NetBackup certificates, the following information is shown: ■ # Certificates. The total number of certificates. Note that the hosts must be online and able to communicate with the NetBackup primary server. ■ Revoked. The number of hosts that have a NetBackup certificate that is revoked. ■ Valid. The number of hosts that have a NetBackup certificate enrolled. ■ Expired. The number of hosts with expired NetBackup certificates. For external certificates, the following information is shown for NetBackup 8.2 and later hosts: ■ # Certificates. The total number of external certificates. Note that the hosts must be online and able to communicate with the NetBackup primary server. ■ Not configured. The number of hosts that do not have an external certificate enrolled. ■ Valid. The number of hosts that have an external certificate enrolled. ■ Expired. The number of hosts with expired external certificates.

Table 3-1 The NetBackup dashboard (*continued*)

Dashboard widget	Description
Tokens	Displays the information about the authorization tokens in your environment. See “Manage NetBackup certificate authorization tokens” on page 427.
Usage reporting	Lists the size of the backup data for the NetBackup primary servers in your organization. This reporting is useful to track capacity licensing. Use the drop-down lists in the top right to select the time period and the view that you want to display. Click on a server name to see specific details for that server. Additional details are available for how to configure NetBackup to display primary server information in this widget. See “Track protected data size on your primary servers” on page 584.
Security events	The Access history view includes a record of logon events. The Audit events view includes the events that users initiate on the NetBackup primary server.

Activity monitor

Use the Activity monitor to monitor and control the following aspects of NetBackup. Updates to the Activity monitor occur as jobs are initiated, updated, and completed.

Jobs	Displays in-process or completed jobs for the primary server. The Jobs tab also displays details about the jobs. See “Job monitoring” on page 49.
Daemons	Displays the status of NetBackup daemons on the primary server. Click Change server to display daemons on a media server in the environment.
Processes	Displays the NetBackup processes that run on the primary server. Click Change server to display processes on a media server in the environment.

Monitor NetBackup daemons

The Activity monitor displays the status of NetBackup daemons on primary and media servers. To start or stop daemons, you must have the applicable RBAC role or similar permissions on the primary or the media server.

Note that not all daemons can be stopped from the NetBackup web UI. On back-level servers, you can still stop and start some services that you cannot in 10.2 and later releases.

To view, stop, or start NetBackup daemons

- 1 On the left, click **Activity monitor**. Then click the **Daemons** tab.
- 2 (Conditional) To manage daemons for a media server in the environment, click **Change server**.
- 3 Locate the daemon.
- 4 On the right, click **Actions**. Then choose from the following actions.

Stop

Stop the selected daemon.

Start

Start the selected daemon.

Monitor NetBackup processes

The Activity monitor displays the status of NetBackup processes on the primary server and the media servers.

To view NetBackup processes

- 1 On the left, click **Activity monitor**. Then click the **Processes** tab.
- 2 (Conditional) To manage processes for a media server in the environment, click **Change server**.

Job monitoring

Use the **Jobs** node in the Activity monitor to monitor the jobs in your NetBackup environment. The default view for jobs is the **List view** that contains the non-hierarchical list of all the jobs. You can also use the **Hierarchical view** to see the hierarchy of parent and child jobs. The role of the parent job is to initiate requested tasks in the form of children jobs.

List view

Hierarchy view

☐	Job ID ↑	Type	Client or display name	Job state
☐	 22322314	Backup	pe...10	Done
☐	 22322315	Backup	pe...10	Done
☐	 22322316	Backup	pe...10	Done
☐	 22322317	Backup	pe...10	Done
☐	 22322318	Backup	pe...10	Done
☐	 22322319	Backup	pe...08	Done

☐	Job ID ↑	Type	Client or display name	Job state
▼ ☐	 22322314	Backup	pe...10	Done
☐	 22322315	Backup	pe...10	Done
☐	 22322316	Backup	pe...10	Done
☐	 22322317	Backup	pe...10	Done
☐	 22322318	Backup	pe...10	Done
▼ ☐	 22322319	Backup	pe...08	Done
☐	 22322320	Backup	pe...08	Done
☐	 22322321	Backup	pe...08	Done
☐	 22322322	Backup	pe...08	Done
☐	 22322323	Backup	pe...08	Done

RBAC permissions for jobs

The type of jobs that you can view and manage depend on the RBAC role that you have. For example, a workload administrator (such as the Default VMware Administrator role) can view and manage only jobs for that workload. In contrast, the Administrator role lets you view and manage all NetBackup jobs.

See [“Workloads that require a custom RBAC role for specific job permissions”](#) on page 50.

Job hierarchy view

If you have an RBAC role that allows access to jobs, you can see the jobs list in the job hierarchy view. For example, the Default VMware Administrator role lets you see VMware jobs in the hierarchy view. However, if you only have access to one or more VMs (asset-level access), no jobs display in the job hierarchy view.

See [“Default RBAC roles”](#) on page 522.

Workloads that require a custom RBAC role for specific job permissions

The NetBackup web UI offers granular job access for certain workloads. This functionality lets you create a custom RBAC role with job permissions for a particular workload.

Note that these workloads do not have a corresponding default RBAC role. When you configure the custom role, the permissions in the **Workloads** card do not apply for these workloads. You can configure job permissions for the following workload types:

BackTrack	Hyper-V	NDMP
DataStore	Informix	PureDisk Export
DB2	Lotus Notes	SAP
Enterprise Vault	SharePoint	Standard
Exchange	MS-Windows	Sybase
FlashBackup	NAS Data Protection	Vault
FlashBackup Windows	NBU Catalog	

To create a custom role with job permissions

- 1 Create a custom RBAC role.
- 2 On the **Assets** tab, locate the workload name and select the job permissions for the workload.

For example, consider that you want to create a custom role so a Hyper-V administrator can view Hyper-V jobs. Locate **Hyper-V** and select the wanted job permissions.

- 3 Select any additional permissions that you want for the role.

For example:

- Other global permissions
- Permissions for protection plans and for credentials

- 4 Add the users you want to assign to the role.

RBAC job permissions for BigData workloads

You cannot configure job permissions specifically for BigData workloads (Hadoop, HBase or MongoDB). To view and manage jobs for BigData, create a role that has the RBAC permissions for all NetBackup jobs.

To configure job permissions

- 1 Create a custom RBAC role.
- 2 Under **Permissions**, click **Assign**.
- 3 On the **Global** tab, expand NetBackup management.
- 4 Locate **Jobs** and select the job permissions you want for the role.
- 5 Add the wanted users to the role.

View a job

For each job that NetBackup runs you can see the following details: the file list and the status of the job, the logged details for the job, and the job hierarchy.

The jobs that you can view depend on the type of RBAC role that you have.

See [“Job monitoring”](#) on page 49.

To view a job and the job details

- 1 On the left, click **Activity monitor**. Then click the **Jobs** tab.
- 2 Click the job name that you want to view.

If you want to open the job in a separate window, at the top right click **Open in new window**.



- 3 On the **Overview** tab you can view information about a job.
 - The **File List** contains the files that are included in the backup image.
 - The **Status** section shows the status and the status codes that are related to the job. Click the status code number to view information about this status code in the Cohesity Knowledge Base.
See the [NetBackup Status Codes Reference Guide](#).
- 4 Click the **Details** tab to view the logged details about a job. You can filter the logs by error type using the drop-down menu.
See [“Search for or filter jobs in the jobs list”](#) on page 54.
- 5 Click the **Job hierarchy** tab to view the complete hierarchy for the job, including any ancestor and any child jobs.
See [“View the jobs in the Hierarchy view”](#) on page 53.

View the jobs in the List view

In the **Jobs** node in the Activity monitor, the list view displays the jobs, without showing the relationship of parent and child jobs.

To view the jobs in the List view

- 1 On the left, click **Activity monitor**. Then click the **Jobs** tab.
- 2 Click the **List view** button.



View the jobs in the Hierarchy view

In the **Jobs** node in the Activity monitor, the hierarchy view displays the jobs so you can see the complete hierarchy of the jobs. This view includes the top-level job (or root job) and its child jobs (if applicable). A child job can, in turn, be a parent of its own child jobs.

To view the jobs in the Hierarchy view

- 1 On the left, click **Activity monitor**. Then click the **Jobs** tab.
- 2 Click the **Hierarchy view** button.



- 3 Locate the top-level job and expand it to see the child jobs.

Jobs: cancel, suspend, restart, resume, delete

Depending on the state of a job, you can perform certain actions on that job.

To manage a job

- 1 On the left, click **Activity monitor**. Then click the **Jobs** tab.
- 2 Select one or more jobs.
- 3 The top menu shows the actions that you can perform for the selected jobs.

Cancel	You can cancel the jobs that have not yet completed. They can be in one of the following states: queued, re-queued, active, incomplete, or suspended. When a parent job is canceled, any child jobs are also canceled.
Suspend	You can suspend backup and restore any jobs that contain checkpoints.
Restart	You can restart the jobs that have completed, failed, or that have been canceled or suspended. A new job ID is created for the new job. Note: Backup Now jobs cannot be restarted.
Resume	You can resume the jobs that have been suspended or are in an incomplete state.
Delete	You can delete the jobs that have completed. When a parent job is deleted, any child jobs are also deleted.

View the logs for a job

In the details for a job, the **Logs** tab displays any logs for the selected job.

To view the logs for a job

- 1 On the left, select **Activity monitor**.
- 2 On the **Jobs** tab, locate and select the link for the job for which you want to view the error logs.
- 3 Select the **Logs** tab.
- 4 You can perform the following tasks:
 - Filter the list of logs by severity. By default, all severity levels display in the list. Select the **Filter** icon then select the severity levels that you want to view: **Critical**, **Error**, **Warning**, or **Information**. The search feature searches within the **Type**, **Process**, and **Description** fields.
 - Copy one or more logs to the Clipboard. Select the check box for one or more logs and select **Copy to clipboard**.
 - View more details for the log. Select the log then select **View details**.

Search for or filter jobs in the jobs list

You can search for jobs in the Activity monitor or create filters to customize the jobs that are displayed.

Search for jobs in the jobs list

The search feature lets you search for the following job information: status code (complete status code #); policy name; client or display name; client; job ID (complete job ID #), or the job's parent ID.

Search for jobs in the jobs list

- 1 On the left, click **Activity monitor**. Then click the **Jobs** tab.
- 2 In the **Search** box, type the keyword you want to find. For example, a client name or a status code number.

Filter the job list

To filter the job list

- 1 On the left, click **Activity monitor**. Then click the **Jobs** tab.
- 2 In the toolbar, click the **Filter** icon.
- 3 Click on a filter that you created. Or, click **All jobs** to display all of the available jobs.

Create a jobs filter

You can create specific filters based on one or more query criteria.

To create a jobs filter

- 1 On the left, click **Activity monitor**. Then click the **Jobs** tab.
- 2 In the toolbar, click the **Filter** icon.
- 3 If you have not created any filters yet, on the left click **Create filter**.
Otherwise, click **Actions > Create**.
- 4 Enter a name and an optional description for the filter.
- 5 Choose if you want the filter to be **Private** or **Public**.

Private

All new filters are private by default. These filters appear in “My list” in the **Manage filters** page. Only the owner can view a private filter.

Public

Public filters are available to all NetBackup users. Any user can view, copy, export, or pin a public filter.

- 6 In the **Query** pane, use the drop-down lists to create a condition.
For example, to view all jobs with the VMware policy type, **Policy type = VMware**.

Query

The screenshot shows a 'Query' pane with a toolbar at the top containing '+ Condition' and '+ Sub-query' buttons. Below the toolbar is a single condition row with three main components: a dropdown menu labeled 'Policy Type', an equals sign '=', and another dropdown menu labeled 'VMware'. To the right of the 'VMware' dropdown is a trash icon for deleting the condition.

- 7 Add any additional conditions for the filter or add a sub-query to apply to a condition.

For example, assume that you want to view all completed jobs that have a status code of 196 or 239. Create the following query:

```
State = Done
AND
  (Status code = 196
   OR
   Status code = 239)
```

Query

AND OR + Condition + Sub-query

State = Done

AND OR + Condition + Sub-query

Status code = 196

Status code = 239

- 8 Choose from the following options:

- To save this query and return to the **Jobs** list, click **Save**.
- To save this query and apply the filter you just created, click **Save and apply**.

Example 1. Query filter for all jobs with the VMware policy type.

Activity monitor

Jobs Daemons Processes Background tasks

Search...

All jobs + VMware

Job ID	Type	Client or display name	Job state	Status code	Schedule	Policy Type	Sche
☐ 68564587	Backup	appliance-10-10-10-10	Done	0	-	VMware	Full t
☐ 68564692	Snapshot	appliance-10-10-10-10	Done	0	-	VMware	Full t
☐ 68564702	Snapshot	appliance-10-10-10-10	Done	0	-	VMware	Full t
☐ 68564707	Snapshot	appliance-10-10-10-10	Done	0	-	VMware	Full t
☐ 68564708	Snapshot	appliance-10-10-10-10	Done	0	-	VMware	Full t

Jobs 394546 (Queued 125, Active 130, Waiting for retry 0, Suspended 0, Incomplete 0, Done 394291) Filter applied (788)

Example 2. Query filter for all jobs that are done and have a status code of 196 or 239.

The screenshot shows the 'Activity monitor' window with the 'Jobs' tab selected. A search filter is applied: 'Code 239 or 196'. The table below lists several failed backup jobs.

Job ID	Type	Client or display name	Job state	Status code	Schedule	Policy Type	Schedule
68879417	Backup	client1	Failed	196	-	Standard	Full bac
68879437	Backup	client1	Failed	196	Full	Standard	Full bac
68879444	Backup	client1	Failed	196	Full	Standard	Full bac
68879445	Backup	client1	Failed	196	Full	Standard	Full bac
68879457	Backup	client1	Failed	196	-	Standard	Full bac
68879482	Backup	client1	Failed	196	Full	Standard	Full bac
68879494	Backup	client1	Failed	196	Full	Standard	Full bac
68585229	Snapshot	client1	Failed	196	-	Hypervisor	Full bac

Jobs 396177 (Queued 55, Active 61, Waiting for retry 0, Suspended 0, Incomplete 0, Done 396061) Filter applied (8)

Edit, copy, or delete a jobs filter

You can edit the query criteria for a jobs filter, copy a filter, or delete a filter that you no longer need.

Edit a jobs filter

To edit a jobs filter

- 1 On the left, click **Activity monitor**. Then click the **Jobs** tab.
- 2 In the toolbar, click **Filter** icon.
- 3 Click **Actions > Manage filters**.
- 4 Click **My list** or **Shared**.

- 5 Select from the following options.

Options with an asterisk (*) are available for any filters that you own.

View	View the details of a filter that you do not own.
Edit*	Edit the filter properties or filter query.
Export	Export the filter to share with another NetBackup user or import the filter into another NetBackup domain.
Make private*	Make a public filter a private filter.
Make public*	Make a private filter a public filter.
Pin	Pin the filter to the jobs filter toolbar.
Delete*	Delete a filter.

- 6 Make the changes that you want to the filter and click **Save**.

Copy a jobs filter

To copy a jobs filter

- 1 On the left, click **Activity monitor**. Then click the **Jobs** tab.
- 2 In the toolbar, click the **Filter** icon.
- 3 Click **Actions > Manage filters**.
- 4 Click **My list** or **Shared**.
- 5 Select the filter that you want to copy.
- 6 Click **View** or **Edit**.
- 7 Make any changes that you want to the filter.
- 8 Click **Copy**.

Delete a jobs filter

To delete a jobs filter

- 1 On the left, click **Activity monitor**. Then click the **Jobs** tab.
- 2 In the toolbar, click the **Filter** icon.
- 3 Click **Actions > Manage filters**.

- 4 Click **My list**.
- 5 Locate the filter that you want to delete and click **Delete > Yes**.

Import or export job filters

The job filter export and import features allow users to share job filters between users or other NetBackup domains.

Import a jobs filter

To import a jobs filter

- 1 On the left, click **Activity monitor**. Then click the **Jobs** tab.
- 2 In the toolbar, click **Filter** icon.
- 3 Click **Actions > Manage filters**.
- 4 Click **My list**.
- 5 Click **Add > Import**.
- 6 Select the filter that you want to import.

Export a jobs filter

To export a jobs filter

- 1 On the left, click **Activity monitor**. Then click the **Jobs** tab.
- 2 In the toolbar, click **Filter** icon.
- 3 Click **Actions > Manage filters**.
- 4 Click **My list** or **Shared**.
- 5 Select the filter that you want to export.
- 6 Click **Export**.

NetBackup exports the filter as a .json file. Note that changing the name of the file does not change the filter name. You can change the filter name after it is imported.

View the status of a redirected restore

A redirected restore may not produce a progress log if the restoring server has no access to write the log files to the requesting server. The name of the requesting server must appear in the server list for the server that performs the restore. (A progress log is an entry on the **Details** tab for a job in the NetBackup web UI.

Progress logs also display in the **View status** dialog box in the **Backup, Archive, and Restore** client interface.)

Consider the following example. `server1` requests a redirected restore from `server2`. To write a log to `server1`, `server1` must appear in the servers list on `server2`.

To add a server requesting a redirected restore to the server list on the restoring server

- 1 In the web UI, sign in to the server that will perform the restore.
For example, sign in to `server2`.
- 2 On the left, select **Hosts > Host properties**.
- 3 Select the primary server.
For example, select `server2`.
- 4 If necessary, click **Connect**. Then click **Edit primary server**.
- 5 Click **Servers**.
- 6 On the **Additional servers** or the **Media servers** tab, click **Add**.
- 7 Enter the name of the server that is requesting the redirected restore.
For example, `server1`.
- 8 Click **Add**.
- 9 Click **Save**.
- 10 Sign into to the requesting server.
For example, `server1`.
Check the **Activity monitor** to determine the success of the restore operation.

Troubleshooting the viewing and managing of jobs

You may see no job results because:

- The keyword or keywords that you searched for do not match any of the details for any jobs.
- You applied a search filter and no jobs match the filter criteria.
- The jobs in the hierarchy view have parent jobs, but you do not have permission to view the parent jobs.
Contact your NetBackup system administrator to get the necessary RBAC role permissions.
- NetBackup limits the number of tabs that you can have open with the Jobs hierarchy view.

If you cannot expand a parent job and see its child jobs, try closing any additional Jobs tabs that you have open.

Some job actions may not be available to workload administrators with limited RBAC permissions on certain assets.

See [“Job actions not available for workload administrators with limited RBAC permissions on assets”](#) on page 61.

Job actions not available for workload administrators with limited RBAC permissions on assets

Note following issues for view and managing jobs with the NetBackup web UI:

- A job does not receive an asset ID until it runs, which means a queued job does not have an asset ID. Users that have roles with more granular asset permissions for a workload are not able to view or cancel queued jobs.
This behavior does not affect users with an RBAC role that has full job permissions or a role that can manage all assets for a particular workload.
- A job does not receive an asset ID if the asset is not yet discovered. Users that have roles with more granular asset permissions for a workload are not able to cancel or restart a job for the asset.
This behavior does not affect users with an RBAC role that has full job permissions or a role that can manage all assets for a particular workload.

Example 1 - VMware administrator with limited asset permissions cannot cancel any queued jobs

Consider a user that has RBAC permissions only for a VMware vCenter or one or more VMs.

- The user cannot see queued jobs for the vCenter or for the VMs.
- Similarly, the user is not able to cancel any queued jobs for the vCenter or for the VMs.

Example 2 - VMware or RHV administrator with limited asset permissions cannot cancel or restart jobs for undiscovered assets

Consider a user that has RBAC permissions only for a VMware vCenter or an RHV server. This user also has one or more job permissions for these assets, but does not have job permissions for all workload assets.

- A new asset is added to the environment, but the discovery process hasn't run yet.
- An existing intelligent group is configured so it includes the new asset.
- When the backup runs, it includes the new asset in the backup.
- The user is not able to cancel or restart a job for the new asset.

Device monitor

This chapter includes the following topics:

- [About the Device Monitor](#)
- [About media mount errors](#)
- [About pending requests and actions](#)

About the Device Monitor

Use the **Device monitor** to manage your tape drives, disk pools, and service requests for operators, as follows:

Media mounts	See “About media mount errors” on page 64.
Pending requests and actions	See “About pending requests and actions” on page 65. See “About pending requests for storage units” on page 66. See “Resubmit a pending request” on page 68. See “Resolve a pending action” on page 67. See “Deny a pending request” on page 68.

Tape drives	See “Change a drive comment” on page 289. See “About downed drives” on page 289. See “Change a drive operating mode” on page 290. See “Clean a tape drive” on page 291. See “Reset a drive ” on page 292. See “Reset the mount time of a drive ” on page 292. See “Set the drive cleaning frequency” on page 293. See “View drive details” on page 293. See “Deny a pending request” on page 68.
Disk pools	More information about disk pools is available in the NetBackup guide for your disk storage option: ■ The <i>NetBackup AdvancedDisk Storage Solutions Guide</i>. ■ The <i>NetBackup Cloud Administrator's Guide</i>. ■ The <i>NetBackup Deduplication Guide</i>. ■ The <i>NetBackup OpenStorage Solutions Guide for Disk</i>. ■ The <i>NetBackup Replication Director Solutions Guide</i>.

About media mount errors

Errors can occur when media is mounted for NetBackup jobs. Depending on the type of error, NetBackup adds the mount request to the pending requests queue or cancels the mount request, as follows:

Adds to the pending requests queue	When NetBackup adds the mount request to the queue, NetBackup creates an operator-pending action. The action appears in the Device monitor. A queued mount request leads to one of the following actions: ■ The mount request is suspended until the condition is resolved. ■ The operator denies the request. ■ The media mount time out is reached.
Cancels the request	When a mount request is automatically canceled, NetBackup tries to select other media to use for backups. (Selection applies only in the case of backup requests.) Many conditions lead to a mount request being automatically canceled instead of queued. When a media mount is canceled, NetBackup selects different media so that the backup is not held up.

When NetBackup selects different media

The following conditions can lead to automatic media reselection:

- The requested media is in a DOWN drive.
- The requested media is misplaced.
- The requested media is write protected.
- The requested media is in a drive not accessible to the media server.
- The requested media is in an offline ACS LSM (Automated Cartridge System Library Storage Module). (ACS robot type only.)
- The requested media has an unreadable barcode. (ACS robot type only.)
- The requested media is in an ACS that is not accessible. (ACS robot type only.)
- The requested media is determined to be unmountable.

About pending requests and actions

In the **NetBackup web UI** click **Storage > Device Monitor**. Then click on the **Device monitor** tab. If requests await action or if NetBackup acts on a request, the request displays in the **Pending requests** pane. For example, if a tape mount requires a specific volume, the request displays in the **Pending requests** pane. If NetBackup requires a specific volume for a restore operation, NetBackup loads or requests the volume.

If NetBackup cannot service a media-specific mount request automatically, it changes the request or action to a pending state.

Table 4-1 Pending states

Pending state	Description
Pending request	Specifies that a pending request is for a tape mount that NetBackup cannot service automatically. Operator assistance is required to complete the request. NetBackup displays the request in the Pending requests pane. NetBackup assigns pending status to a mount request when it cannot determine the following: ■ Which standalone drive to use for a job. ■ Which drive in a robot is in Automatic Volume Recognition (AVR) mode.

Table 4-1 Pending states (*continued*)

Pending state	Description
Pending action	Specifies that a tape mount request becomes a pending action when the mount operation encounters problems, and the tape cannot be mounted. Operator assistance is required to complete the request, and NetBackup displays an action request in the Pending requests pane. Pending actions usually occur with drives in robotic libraries.

About pending requests for storage units

In the **NetBackup web UI**, click **Storage > Device Monitor**. Then click on the **Device monitor** tab.

The following tape mount requests do not appear in the **Pending requests** pane:

- Requests for backups
- Requests for a tape that is required as the target of a duplication operation

These requests are for resources in a storage unit and therefore are not for a specific volume. NetBackup does not assign a mount request for one storage unit to the drives of another storage unit automatically. Also, you cannot reassign the mount request to another storage unit.

If the storage unit is not available, NetBackup tries to select another storage unit that has a working robot. If NetBackup cannot find a storage unit for the job, NetBackup queues the job (a **Queued** state appears in the Activity monitor).

You can configure NetBackup so that storage unit mount requests are displayed in the **Device monitor** if the robot or drive is down. Pending requests display in the **Device monitor**, and you can assign these mount requests to drives manually.

Resolve a pending request

Use the following procedure to resolve a pending request.

To resolve a pending request

- 1 Insert the requested volume in a drive that matches the density of the volume that was requested.
- 2 Open the NetBackup web UI.
- 3 On the left, click **Storage > Tape storage**. Then click on the **Device monitor** tab.
- 4 In the **Pending requests** pane, select the request and note the contents of the following columns of the request:

- Density
 - Recorded media ID
 - Mode
- 5 Find a drive type that matches the density for the pending request.
 - 6 Verify that the drive is up and not assigned to another request.
 - 7 Locate the drive. Then ensure that the drive and the pending request are on the same host.
 - 8 If necessary, get the media, write-enable it, and insert it into the drive.
 - 9 Wait for the drive to become ready, as explained in the vendor's drive equipment manual.
 - 10 Locate the request. Then click **Actions > Assign request**.
 - 11 Verify that the request was removed from the **Pending requests** pane.
 - 12 Click on the drive name, then click on the **Drive status** tab.
 Verify that the job request ID appears in the Request ID column for the drive.

Resolve a pending action

A pending action is similar to a pending request. For a pending action, NetBackup determines the cause of the problem and issues an instruction to the operator to resolve the problem.

Use the following procedure to resolve a pending action.

To resolve a pending action

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**. Then click on the **Device monitor** tab.
- 3 In the **Pending requests** pane, locate the pending action.
- 4 Click **Actions > Display pending action**.
- 5 Review the list of possible actions and click **OK**.
- 6 Correct the error condition and either resubmit the request or deny the request.

See ["Resubmit a pending request"](#) on page 68.

See ["Deny a pending request"](#) on page 68.

Resubmit a pending request

After you correct a problem with a pending action, you can resubmit the request.

If the problem is a volume missing from a robot, first locate the volume, insert it into the robot, and then update the volume configuration. Usually, a missing volume was removed from a robot and then requested by NetBackup.

To resubmit a request

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**. Then click on the **Device monitor** tab.
- 3 In the **Pending requests** pane, locate the request.
- 4 Click **Actions > Resubmit request**.

Deny a pending request

Some situations may require that you deny requests for service. For example, when a drive is not available, you cannot find the volume, or the user is not authorized to use the volume. When you deny a request, NetBackup sends an appropriate status message to the user.

To deny a request

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**. Then click on the **Device monitor** tab.
- 3 In the **Pending requests** pane, locate the request.
- 4 Then click **Actions > Deny request**.

Notifications

This chapter includes the following topics:

- [Job notifications](#)
- [NetBackup event notifications](#)

Job notifications

The following types of email notifications are available for NetBackup jobs.

- Notifications when job failures occur. NetBackup supports the ticketing systems that use inbound email service for ticket creation.
See [“Send email notifications for job failures”](#) on page 69.
- Notifications to the backup administrator about backups with a non-zero status.
See [“Send notifications to the backup administrator about failed backups”](#) on page 72.
- Notifications to the host administrator about successful and failed backups for a specific host.
See [“Send notifications to a host administrator about backups”](#) on page 73.

Send email notifications for job failures

You can configure NetBackup to send email notifications when job failures occur. This way administrators spend less time monitoring NetBackup for job failures and manually creating tickets to track issues. NetBackup supports the ticketing systems that use inbound email service for ticket creation.

See [“Status codes that generate alerts”](#) on page 71.

NetBackup generates alerts based on certain job failure conditions or NetBackup status codes. Alerts that are similar or have a similar reason for failure are marked as duplicates. Email notifications for duplicate alerts are not sent for the next 24

hours. If a notification cannot be sent, NetBackup retries every 2 hours, up to three attempts.

NetBackup audits an event if changes are made to the alert settings or when it cannot generate an alert or send an email notification.

See [“About NetBackup auditing”](#) on page 413.

Prerequisites

Review the following requirements before you configure email notifications using a ticketing system.

- The ticketing system is up and running.
- The SMTP server is up and running.
- A policy is configured in the ticketing system to create tickets (or incidents) based on the inbound emails that NetBackup sends.

To configure email notifications

- 1 At the top right, click **Settings > Email notifications**.
- 2 Go to the **Email notifications** tab.
- 3 Turn on **Send email notifications**.
- 4 Enter the email information including the recipient's email address, the sender's email address, and the email sender's name.
- 5 Enter the SMTP server details including the SMTP server name and port number.

Provide the SMTP username and password if you have specified the credentials earlier on the SMTP server.
- 6 Select **Save**.
- 7 Log on to the ticketing system to view the tickets that were created based on NetBackup alerts.

Exclude specific status codes from email notifications

You can exclude specific status codes so that email notifications are not sent for these errors.

To exclude specific status codes

- 1 At the top right, click **Settings > Email notifications**.
- 2 Select the **Excluded notifications** tab.
- 3 Go to **Job failures**.

- 4 If necessary, clear **Do not send any notifications**.
- 5 Enter the status codes or a range of status codes (separated by commas) for which you do not want to receive email notifications.
- 6 Select **Save**.

Sample email notification for an alert

An email notification for an alert contains information about the primary server, job, policy, schedule, and error. Emails may contain other information based on the type of job. For example, for VMware job failures, details such as vCenter Server and ESX host are present in the email notification.

Example email notification:

Primary Server: primary1.example.com

Client Name: client1.example.com

Job ID: 50

Job Start Time: 2018-05-17 14:43:52.0

Job End Time: 2018-05-17 15:01:27.0

Job Type: BACKUP

Parent Job ID: 49

Policy Name: Win_policy

Policy Type: WINDOWS_NT

Schedule Name: schedule1

Schedule Type: FULL

Status Code: 2074

Error Message: Disk volume is down

Status codes that generate alerts

The NetBackup web UI supports alerts for VMware job failures and retains the alerts for 90 days. NetBackup generates alerts for the supported status codes for following job types: backup, snapshot, snapshot replication, index from snapshot, and backup from snapshot. For the complete list of status codes for which alerts are generated, refer to the information for alert notification status codes in the [NetBackup Status Codes Reference Guide](#).

[Table 5-1](#) lists some of the conditions or status codes for which alerts are generated. These alerts are sent to the ticketing system through email notifications.

Table 5-1 Examples of status codes that generate alerts

Status code	Error message
10	Allocation failed
196	Client backup was not attempted because backup window closed
213	No storage units available for use
219	The required storage unit is unavailable
2001	No drives are available
2074	Disk volume is down
2505	Unable to connect to the database
4200	Operation failed: Unable to acquire snapshot lock
5449	The script is not approved for execution
7625	SSL socket connection failed

Send notifications to the backup administrator about failed backups

You can send notifications to the backup administrator about backups with a non-zero status.

On UNIX, NetBackup uses the mail transfer agent sendmail to send email notifications. For Windows, NetBackup requires that an application to transfer messages using SMTP is installed and that the `nbmail.cmd` script is configured on the Windows hosts that send notifications.

See [“Configure the nbmail.cmd script on the Windows hosts”](#) on page 73.

To configure notifications for the backup administrator of a NetBackup host, see the following topic.

See [“Send notifications to a host administrator about backups”](#) on page 73.

To send notifications to the backup administrator about failed backups

- 1 On the left, select **Hosts > Host properties**.
- 2 Select the primary server.
- 3 If necessary click **Connect**. Then click **Edit primary server**.

- 4 Click **Global attributes**.
- 5 Enter the email address of the administrator. (Separate multiple addresses with commas.)
- 6 Click **Save**.

Send notifications to a host administrator about backups

You can send notifications to the host administrator about successful and failed backups for a specific host.

On UNIX, NetBackup uses the mail transfer agent sendmail to send email notifications. Windows requires that an application to transfer messages with SMTP is installed. You also must configure the `nbmail.cmd` script on the Windows hosts that send notifications.

See [“Configure the nbmail.cmd script on the Windows hosts”](#) on page 73.

To send notifications for backups of a specific host

- 1 On the left, select **Hosts > Host properties**.
- 2 Select the client.
- 3 If necessary click **Connect**. Then click **Edit client**.
- 4 Click **Universal settings**.
- 5 Choose how to send the email notifications.
 - To send email notifications from the client, select **Client sends email**.
 - To send email notifications from the server, select **Server sends email**.
- 6 Enter the email address of the host administrator. (Separate multiple addresses with commas.)
- 7 Click **Save**.

Configure the nbmail.cmd script on the Windows hosts

For Windows hosts to send and receive email notifications about backups, the `nbmail.cmd` script must be configured on the applicable hosts.

To configure the nbmail.cmd script on the Windows hosts

- 1 Create a backup copy of `nbmail.cmd`.
- 2 On the primary server, locate the following script:


```
install_path\NetBackup\bin\goodies\nbmail.cmd
```
- 3 Copy the script to the following directory on the applicable hosts:

`install_path\NetBackup\bin\`

Primary and media server NetBackup sends notifications from the server if you configure the following setting:

- The **Administrator's email address** in Global Attributes.
- The **Server sends email** option in the **Universal Settings**.

Client. NetBackup sends notifications from the client if you configure the following setting:

- The **Client sends email** option in the **Universal Settings**.

4 Use a text editor to open `nbmail.cmd`.

The following options are used in the script:

<code>-s</code>	The subject line of the email
<code>-t</code>	Indicates who receives the email.
<code>-i</code>	The originator of the email, though it is not necessarily known to the mail server. The default (<code>-i Netbackup</code>) shows that the email is from NetBackup.
<code>-server</code>	The name of the SMTP server that is configured to accept and relay emails.
<code>-q</code>	Suppresses all output to the screen.

5 Adjust the lines as follows:

- Remove `@REM` from each of the five lines to activate the necessary sections for BLAT to run.
- Replace `SERVER_1` with the name of the mail server. For example:

```
@IF "%~4"==" " (
blat %3 -s %2 -t %1 -i Netbackup -server emailserver.company.com -q
) ELSE (
blat %3 -s %2 -t %1 -i Netbackup -server emailserver.company.com -q -attach %4
)
```

6 Save `nbmail.cmd`.

NetBackup event notifications

To make NetBackup administrators aware of important system events, NetBackup regularly queries system logs and displays notifications about the events.

Note: Job events are not included with these notifications. See job details in the **Activity Monitor** for information about job events.

A **Notifications** icon is located at the top right in the web UI. You can click the icon to open the **Notifications** window and view a list of critical notifications 10 at a time. If a number is displayed with the icon, it indicates how many unseen critical messages exist. After you have opened the window, the number is reset.

From the window, you can choose to see a more comprehensive list of all notifications. Each event has a category for its NetBackup or external component and is assigned a severity level:

- Error
- Critical
- Warning
- Information
- Debug
- Notice

You can sort, filter, and search the list. The comprehensive list also lets you review details about each event. The details include the full description as well as any appropriate extended attributes.

NetBackup notifications are not available if the NetBackup Messaging Broker (`nbmqbroker`) is not running. See the *NetBackup Troubleshooting Guide* for information about restarting the service.

View notifications

To view notifications

- 1 At the top right, click the **Notifications** icon to view a list of critical notifications 10 at a time.

Note: If a number is displayed with the icon, it indicates how many unseen critical messages exist. After you have opened the **Notifications** window, the number is reset.

Click **Load 10 more** to view the next 10 notifications. After you have viewed 30 notifications, click **Show all** to view any remaining messages.

Use **Refresh** to load the most recent notifications again.

- 2 To view all notifications, click **Show all** to open the **Events** page. On the page, you can do the following:
 - Click an event to view its details. The details include the full description as well as extended attributes.
 - To sort the list, click any of the column headings except **Description**. Events are sorted by default by the date received.
 - To filter events, click **Filter**. You can filter by **Severity** and **Timeframe**. In the **Filters** menu, select the parameter values you want to filter by, and then click **Apply filters**.
To remove all filters, click **Clear all**.
 - To search for events, enter the search string in the **Search** field. You can search for values in all columns except **Description** and **Received**.

Modify or disable NetBackup event notifications in the web UI

You can disable specific types of NetBackup event notifications that appear in the web UI, or modify their severity and priority, by making changes to the `eventlog.properties` file on the NetBackup primary server:

- Windows:
`install_path\var\global\wmc\h2Stores\notifications\properties`
- UNIX:
`/usr/opensv/var/global/wmc/h2Stores/notifications/properties`

Disable event notifications

To disable event notifications

- ◆ Add a `DISABLE` entry in the `eventlog.properties` file in one of the following formats:

```
DISABLE.NotificationType = true
```

```
Or DISABLE.NotificationType.Action = true
```

```
Or DISABLE.namespace
```

For valid *NotificationType* and *Action* values, see the following topic.

See [“NetBackup event types supported with notifications”](#) on page 78.

For example:

- To disable notifications about all storage unit events:

```
DISABLE.StorageUnit = true
```

- To disable only notifications about create storage unit events:

```
DISABLE.StorageUnit.CREATE = true
```

- To disable only notifications about update to storage unit events using a namespace:

```
DISABLE.eventlog.vrts.nbu.emm.storageunit.update = true
```

Modify event notifications

To modify the priority or severity of event notifications

- ◆ Add or change an entry in the `eventlog.properties` file in one of the following formats:

```
NotificationType.Action.priority = value
```

```
Or NotificationType.Action.severity = value
```

Valid priority values are: LOW, MEDIUM, HIGH

Valid severity values are: CRITICAL, ERROR, WARNING, INFO, DEBUG

For example:

- To set priority and severity for create storage unit events:

```
StorageUnit.CREATE.priority = LOW
```

```
StorageUnit.CREATE.severity = INFO
```

Note: It can take up to one minute for the events of type Policy, SLP, and Catalog to generate after the corresponding action has been performed.

NetBackup event types supported with notifications

The following NetBackup event types support event notifications in the NetBackup web UI.

Table 5-2 NetBackup event types supported with notifications

Event type and notification type value	Action	Severity	Sample notification message
Autodiscovery and Discover Now <code>AutoDiscoveryEvent</code>	no actions	INFO	An appropriate notification is generated when an autodiscovery action or a Discover Now action is performed for VMWare, RHV, or Cloud servers.
	no actions	CRITICAL	Note: An appropriate notification is generated when an autodiscovery action or a Discover Now action fails for VMWare, RHV, Nutanix, or Cloud servers. Note: An appropriate notification is generated when an autodiscovery action or a Discover Now action fails for VMWare, RHV, or Cloud servers.
CRL Health	Not applicable	CRITICAL	The CRL on host \$ {hostName} is not refreshed.
Catalog Backup Health	Not applicable	CRITICAL	One or more users who can access the identity files that need to be backed up as part of the disaster recovery (DR) package, do not exist on the system.
Catalog Image Expiration <code>Catalog</code> Note: Also applicable for manual image expiration.	Not applicable	CRITICAL	Event for Catalog Image received. No additional details found. Catalog Image <i>Image_Name</i> was modified. Catalog Image <i>Image_Name</i> expired.
cDOT Client <code>cDOTClientEvent</code>	CREATE	INFO	<i>{Cluster_Data_ONTAP_Client_Name}</i> was added as a cDOT client.
	DELETE	CRITICAL	<i>{Cluster_Data_ONTAP_Client_Name}</i> was deleted as a cDOT client.
Certificate Health	Not applicable	CRITICAL	The certificate for host \$ {hostName} is going to expire soon.

Table 5-2 NetBackup event types supported with notifications (*continued*)

Event type and notification type value	Action	Severity	Sample notification message
Client <i>ClientEvent</i>	CREATE	INFO	The client { <i>Client_Name</i> } was created.
	DELETE	CRITICAL	The client { <i>Client_Name</i> } was deleted.
	UPDATE	INFO	The client { <i>Client_Name</i> } was updated.
NetBackup Configuration Health	Not applicable	CRITICAL	The NetBackup configuration file contains multiple <i>CLIENT_NAME</i> entries.
NetBackup Configuration Health	Not applicable	CRITICAL	The service user does not have the required permissions on one or more links or junction target directories. Run the '<i>Install_Path\NetBackup\bin\goodies\nbserviceusercmd.exe -addAcl</i> ' command to assign the correct permissions. The service user does not have the required permissions on one or more soft link target directories. The service user does not have the required permissions on ALTPATH directories that are configured for one or more clients. Run the '<i>Install_Path\NetBackup\bin\goodies\nbserviceusercmd.exe -addAcl</i> ' command to assign the correct permissions.
NetBackup Configuration Health	Not applicable	INFO	Assigned the execute permission to the service user on one or more NetBackup directories.
NetBackup Configuration Health	Not applicable	WARNING	Could not assign the execute permission to the service user on one or more NetBackup directories.
DBPaaS Operation RCA	Not applicable	CRITICAL	Cannot complete backup. See the Root Cause Identifier (RCA) link for more information.
Drive <i>DriveChange</i>	CREATE	INFO	The drive { <i>Drive_Name</i> } was created for host { <i>Host_Name</i> }.
	DELETE	CRITICAL	The drive { <i>Drive_Name</i> } was deleted for host { <i>Host_Name</i> }.
	UPDATE	INFO	The drive {<i>Drive_Name</i>} was updated for host {<i>Host_Name</i>}. Note: A notification message like this one is generated when a drive is updated for a particular host or when a drive state is changed to UP or DOWN.

Table 5-2 NetBackup event types supported with notifications (*continued*)

Event type and notification type value	Action	Severity	Sample notification message
Isilon Client <code>IsilonClientEvent</code>	CREATE	INFO	<i>{Isilon_Filer_Client_Name}</i> was added as an Isilon client.
	DELETE	CRITICAL	<i>{Isilon_Filer_Client_Name}</i> was deleted as an Isilon client.
KMS Certificate Expiration <code>KMSCredentialStatus</code>	EXPIRY	WARNING	The certificate that is used to communicate with the KMS server <i>{KMS_Server_Name}</i> is about to expire in <i>{days_to_expiration}</i> . If the certificate is not renewed on time, communication with the KMS server fails.
Library Event - Robot <code>Library</code>	CREATE	INFO	The library <i>{Library_Name}</i> was created for host <i>{Host_Name}</i> .
	DELETE	CRITICAL	The library <i>{Library_Name}</i> was deleted for host <i>{Host_Name}</i> .
	UPDATE	INFO	The library <i>{Library_Name}</i> was updated for host <i>{Host_Name}</i> .
Machine [Primary/Media/Cluster] <code>Machine</code>	CREATE	INFO	The host <i>{Host_Name}</i> was created.
	DELETE	CRITICAL	The host <i>{Host_Name}</i> was deleted.
Media <code>Media</code>	CREATE	INFO	The media <i>{Media_ID}</i> was created.
	DELETE	CRITICAL	The media <i>{Media_ID}</i> was deleted.
	UPDATE	INFO	The media <i>{Media_ID}</i> was updated.
Media Group <code>MediaGroup</code>	CREATE	INFO	The media group <i>{Media_Group_ID}</i> was created.
	DELETE	CRITICAL	The media group <i>{Media_Group_ID}</i> was deleted.
	UPDATE	INFO	The media group <i>{Media_Group_ID}</i> was updated.
Media Pool <code>MediaPool</code>	CREATE	INFO	The media pool <i>{Media_Pool_ID}</i> was created.
	DELETE	CRITICAL	The media pool <i>{Media_Pool_ID}</i> was deleted.

Table 5-2 NetBackup event types supported with notifications (*continued*)

Event type and notification type value	Action	Severity	Sample notification message
	UPDATE	INFO	The media pool <i>{Media_Pool_ID}</i> was updated.
Message Broker Service Status <i>ServiceStatus</i>	RUNNING	INFO	The NetBackup Messaging Broker service is running. NetBackup internal notifications are now enabled.
	STOPPED	INFO	The NetBackup Messaging Broker service is stopped. NetBackup internal notifications are now disabled.
Policy <i>Policy</i> Note: When possible, an aggregated policy event for two or more policy actions is created.	Create	INFO	The policy <i>{Policy_Name}</i> was created. Event for Policy received. No additional details found.
	Update	INFO or CRITICAL	The policy <i>{Policy_Name}</i> was activated. The policy <i>{Policy_Name}</i> was deactivated. The policy <i>{Policy_Name}</i> was updated. The client <i>{Policy_Name}</i> was added to the policy <i>\${policyName}</i> . The client <i>{Policy_Name}</i> was removed from the policy <i>{Policy_Name}</i> . The schedule <i>{Policy_Name}</i> was added to the policy <i>\${Policy_Name}</i> . The schedule <i>{Policy_Name}</i> was removed from the policy <i>{Policy_Name}</i> .
	Delete	CRITICAL	The policy <i>{Policy_Name}</i> was deleted.
Protection Plan <i>ProtectionPlan</i>	Create	INFO	Received an event for protection plan. The protection plan <i>Protection_Plan_Name</i> is created. The protection plan <i>Protection_Plan_Name</i> is created from existing NetBackup policy.
	Update	INFO	The protection plan <i>Protection_Plan_Name</i> is updated.
	Delete	CRITICAL	The protection plan <i>Protection_Plan_Name</i> is deleted.

Table 5-2 NetBackup event types supported with notifications (*continued*)

Event type and notification type value	Action	Severity	Sample notification message
Protection Plan Subscription <i>ProtectionPlanSubscription</i>	Create	INFO	Received an event for protection plan subscription. The <i>Asset_Class Asset_Display_Name</i> is subscribed to protection plan <i>Protection_Plan_Name</i> .
	Update	INFO	The <i>Asset_Class Asset_Display_Name</i> subscription with protection plan <i>Protection_Plan_Name</i> is updated.
	Delete	CRITICAL	The <i>Asset_Class Asset_Display_Name</i> is unsubscribed from protection plan <i>Protection_Plan_Name</i> .
Retention Event <i>RetentionEvent</i>	UPDATE	INFO	Retention level has been changed.
Storage life cycle policy <i>SLP</i>	Create	INFO	Event for Storage Lifecycle Policy received. No additional details found. The Storage Lifecycle Policy <i>{Policy_Name}</i> was created.
	Delete	CRITICAL	The Storage Lifecycle Policy <i>{Policy_Name}</i> was deleted. The Storage Lifecycle Policy <i>{Policy_Name}</i> with version <i>Version_Number</i> was deleted.
Storage life cycle policy state change <i>SlpVersionActInactEvent</i>	UPDATE	INFO	The SLP version <i>{Version}</i> was changed.
Storage Unit <i>StorageUnit</i> Note: Any change to a basic disk staging schedule (DSSU), such as adding, deleting, or modifying, generates relevant storage unit notifications. With those notifications, some additional policy notifications are also generated with policy name <i>__DSSU_POLICY_{Storage_Unit_Name}</i> .	CREATE	INFO	The storage unit <i>{Storage_Unit_Name}</i> was created.
	DELETE	CRITICAL	The storage unit <i>{Storage_Unit_Name}</i> was deleted.
	UPDATE	INFO	The storage unit <i>{Storage_Unit_Name}</i> was updated.

Table 5-2 NetBackup event types supported with notifications (*continued*)

Event type and notification type value	Action	Severity	Sample notification message
Storage Unit Group <code>StorageUnitGroup</code>	CREATE	INFO	The storage unit group { <i>Storage_Unit_Group_Name</i> } was created.
	DELETE	CRITICAL	The storage unit group { <i>Storage_Unit_Group_Name</i> } was deleted.
	UPDATE	INFO	The storage unit group { <i>Storage_Unit_Group_Name</i> } was updated.
	UPDATE	INFO	The storage service { <i>Storage_Service_Name</i> } was updated.
Usage Reporting <code>UsageReportingEvent</code>	No actions	INFO or ERROR	The usage report generation has started. The usage report is generated successfully. Failed to generate the usage report. For more details, refer to the gather and report logs in the parent directory.
VMware Discovery <code>TAGSDISCOVERYEVENT</code>	no actions	INFO	VMware tags cannot be retrieved.
Web Truststore Health	Not applicable	CRITICAL	One or more files and / or directories do not have appropriate web service user permissions.

About configuring automatic notification cleanup tasks

By default, NetBackup runs event notification cleanup tasks every 4 hours. Up to 10,000 event records are stored for up to 3 days in the event database. During the cleanup tasks, NetBackup removes the older notifications from the database.

You can change how often the cleanup tasks run, how many event records are kept at one time, and the number of days a record is retained.

From a command line, use `bpsetconfig` or `bpgetconfig` to change the parameter values listed in [Table 5-3](#). See the *NetBackup Command Reference Guide* for more information about these commands.

You can also change the parameter values with the following APIs:

- `GET/config/hosts/{hostId}/configurations`
- `POST/config/hosts/{hostId}/configurations`
- `GET/config/hosts/{hostId}/configurations/configurationName` (for a specific property)

- PUT/config/hosts/{hostId}/configurations/configurationName
- DELETE/config/hosts/{hostId}/configurations/configurationName

See the *NetBackup 11.0 API Reference* on [SORT](#) for more information about these APIs.

Table 5-3 Configurable parameters for automatic notification cleanup tasks

Parameter and description	Minimum value	Default value	Maximum value
EVENT_LOG_NOTIFICATIONS_COUNT The maximum number of records that are stored, after which the cleanup process removes the oldest record, overriding the retention value.	1000	10000	100000
EVENT_LOG_NOTIFICATIONS_RETENTION_IN_HOURS The number of hours for which the events are stored in the database.	24 (hours)	72 (hours)	168 (hours)
EVENT_LOG_NOTIFICATIONS_CLEANUP_INTERVAL_IN_HOURS The frequency at which the event cleanup service runs.	1 (hour)	4 (hours)	24 (hours)

Registering the data collector

This chapter includes the following topics:

- [About the data collector](#)
- [Register the data collector with Cohesity Alta View](#)
- [Renew Cohesity Alta View token](#)
- [Register the data collector with NetBackup IT Analytics](#)
- [View and modify the data collector registration](#)
- [Unregister the data collector](#)

About the data collector

The data collector collects metadata from NetBackup and sends information such as policies, jobs, image records to Alta View or NetBackup IT Analytics. Based on the information that the data collector has sent, these applications monitor, manage, and report on NetBackup domains.

See [“Register the data collector with Cohesity Alta View”](#) on page 86.

See [“Register the data collector with NetBackup IT Analytics”](#) on page 87.

To receive data from the data collector, you need to register Cohesity Alta View or NetBackup IT Analytics with the data collector.

Note: Either Cohesity Alta View or NetBackup IT Analytics can be registered with a single data collector at a time.

Register the data collector with Cohesity Alta View

Cohesity Alta View is a centralized management platform to manage multiple NetBackup domains. It provides global enterprise data protection visibility and operations. It is a cloud-based management console unifying the protection and management of on-premises and cloud workloads from a single interface, delivering simplified policy management, centralized visibility, and flexible protection strategies.

For more information, refer to the *Cohesity Alta View Help*.

To enable Cohesity Alta View to collect data from NetBackup, you must register the data collector that you have on your primary server with Cohesity Alta View using the NetBackup web UI.

To register the data collector with Cohesity Alta View

- 1 On the top right, click **Settings > Data collector registration**.
- 2 Click **Register with Cohesity Alta View**.
- 3 Click **Choose file** to select the registration file (JSON) that you have downloaded using the Cohesity Alta View UI earlier.

See the 'Complete domain registration for NetBackup 10.1.1 and later' topic in the *Cohesity Alta View Help*.

- 4 Select the **Use proxy server** option and specify the proxy server settings.
This is an optional step.

- 5 Click **Register**.

After the registration with the data collector, you can monitor, manage, and report on NetBackup domains using the Cohesity Alta View UI and the Cohesity Alta View Reports UI.

After the registration, you can access Cohesity Alta View using the NetBackup web UI. The Cohesity Alta View option is added on the left pane in the UI.

Renew Cohesity Alta View token

After the data collector is registered with Cohesity Alta View for data collection, the connection between the Cohesity Alta View server and the NetBackup primary server is established.

You can view the registration and the connection status using the NetBackup web UI.

See [“View and modify the data collector registration”](#) on page 88.

However, the Cohesity Alta View server may be disconnected from the primary server if the token of Cohesity Alta View has expired.

To renew the Cohesity Alta View token

- 1 On the top right, click **Settings > Data collector registration**.
- 2 Verify if the **WebSocket status** is disconnected because the token has expired.
- 3 In the Cohesity Alta View UI, on the **NetBackup domains > Hosts** tab, select the primary server that is disconnected from this Cohesity Alta View server.
- 4 Click **Actions > Generate token**.
Copy the token.
- 5 In the NetBackup web UI, on the **Data collector registration** screen, click **Renew Cohesity Alta View token**.
- 6 In the **Renew Cohesity Alta View token** dialog box, enter the token that you have generated in the Cohesity Alta View UI.
- 7 Click **Renew**.

Register the data collector with NetBackup IT Analytics

NetBackup IT Analytics is the storage resource management platform that enables IT organizations to integrate storage and backup solutions to address rapid growth and declining budgets.

For more information, see the *NetBackup IT Analytics User Guide*.

To enable NetBackup IT Analytics to collect data from NetBackup, you must register the data collector that you have on your primary server with NetBackup IT Analytics using the NetBackup web UI.

If NetBackup IT Analytics portal is hosted on on-premise, you must register the data collector with the portal.

To register the data collector with NetBackup IT Analytics

- 1 On the top right, click **Settings > Data collector registration**.
- 2 Click **Register with NetBackup IT Analytics**.
- 3 Click **Choose file** to select the registration file (JSON) that you have downloaded using the NetBackup IT Analytics portal earlier.

See the 'Add/Edit Data Collectors' topic in the *NetBackup IT Analytics User Guide*.

- 4 Select the **Use proxy server** option and specify the proxy server settings.

This is an optional step.

- 5 Click **Register**.

After the registration with the data collector, you can monitor, manage, and report on NetBackup domains using NetBackup IT Analytics.

View and modify the data collector registration

The data collector collects metadata from NetBackup and sends it to Veritas Alta View or NetBackup IT Analytics. To receive data from the data collector, you need to register Veritas Alta View or NetBackup IT Analytics with the data collector.

Note: Either Cohesity Alta View or NetBackup IT Analytics can be registered with a single data collector at a time.

See [“About the data collector”](#) on page 85.

See [“Register the data collector with Cohesity Alta View”](#) on page 86.

See [“Register the data collector with NetBackup IT Analytics”](#) on page 87.

After the data collector is registered with Veritas Alta View or NetBackup IT Analytics, you can view the registration and the connection status using the NetBackup web UI.

You can also modify some of the registration parameters.

To view and modify the data collector registration

- ◆ On the top right, click **Settings > Data collector registration**.

NetBackup indicates if the data collector on the NetBackup primary server is registered with Cohesity Alta View or NetBackup IT Analytics.

If the data collector on the NetBackup primary server is registered with Cohesity Alta View, the following details are displayed.

- **Proxy server** - Shows if the proxy server is enabled or disabled.
Click **Edit** to modify the proxy server settings.
- **Data collection** - Shows if the data collection is enabled or disabled.
Turn on the option if you want the data collector to start collecting data from the NetBackup primary server.
- **WebSocket status** - Shows the status of the connection between the data collector and the Cohesity Alta View server.
The WebSocket may be disconnected in certain cases.

For example, the data collector is disconnected from the Cohesity Alta View server after the Cohesity Alta View token has expired.

See [“Renew Cohesity Alta View token”](#) on page 86.

If the data collector on the NetBackup primary server is registered with NetBackup IT Analytics, the following details are displayed.

- **Proxy server** - Shows if the proxy server is enabled or disabled.
Click **Edit** to modify the proxy server settings.
- **Data collection** - Shows if the data collection is enabled or disabled.
Turn on the option if you want the data collector to start collecting data from the NetBackup primary server.

Unregister the data collector

To stop collecting data from NetBackup, you must unregister the data collector that you have registered earlier with Cohesity Alta View or NetBackup IT Analytics.

If you want to change registration from Cohesity Alta View to NetBackup IT Analytics portal or from NetBackup IT Analytics portal to Cohesity Alta View, you must first unregister the existing configuration.

To unregister the data collector

- 1 On the top right, click **Settings > Data collector registration**.
- 2 Click **Unregister data collector**.

Configuring hosts

- [Chapter 7. Managing host properties](#)
- [Chapter 8. Managing credentials for workloads and systems that NetBackup accesses](#)
- [Chapter 9. Managing deployment](#)

Managing host properties

This chapter includes the following topics:

- [Overview of host properties](#)
- [View or edit the host properties of a server or client](#)
- [Host information and settings in Host properties](#)
- [Reset a host's attributes](#)
- [Active Directory properties](#)
- [Backup pool host properties](#)
- [Busy file settings properties](#)
- [Clean up properties](#)
- [Client name properties](#)
- [Client attributes properties](#)
- [Client settings properties for UNIX clients](#)
- [Client settings properties for Windows clients](#)
- [Cloud Storage properties](#)
- [Credential access properties](#)
- [Data Classification properties](#)
- [Default job priorities properties](#)
- [Distributed application restore mapping properties](#)
- [Encryption properties](#)

- Enterprise Vault properties
- Enterprise Vault hosts properties
- Exchange properties
- Exclude list properties
- Fibre transport properties
- Firewall properties
- General server properties
- Global attributes properties
- Logging properties
- Lotus Notes properties
- Media properties
- Network properties
- Network settings properties
- Nutanix AHV access hosts
- Port ranges properties
- Preferred network properties
- Properties setting in host properties
- RHV access hosts properties
- Resilient network properties
- Resource limit properties
- Restore failover properties
- Retention periods properties
- Scalable Storage properties
- Servers properties
- SharePoint properties
- SLP settings properties
- Throttle bandwidth properties

- [Timeouts properties](#)
- [Universal settings properties](#)
- [UNIX client properties](#)
- [UNIX Server properties](#)
- [User account settings properties](#)
- [VMware access hosts properties](#)
- [Windows client properties](#)
- [Configuration options not found in the host properties](#)
- [About using commands to change the configuration options on UNIX or Linux clients and servers](#)

Overview of host properties

The configuration options within the **Host properties** let an administrator customize NetBackup to meet specific site preferences and requirements.

To change the properties of another client or server, the NetBackup server that you signed in to must be in the **Servers** list on the other system.

For example, if you logged on to *server_1* and want to change a setting on *client_2*, *client_2* must include *server_1* in its **Servers** list.

See “[Servers properties](#)” on page 197.

A NetBackup administrator can use one of the following methods to read or set the default configuration options. Some options cannot be configured by using the **NetBackup web UI**.

Table 7-1 NetBackup Host properties configuration methods

Method	Description
NetBackup Web UI interface	Most properties are listed in the NetBackup web UI in Hosts > Host properties . Depending on the host to be configured, select the Primary server , Media server , or Client .

Table 7-1 NetBackup Host properties configuration methods (*continued*)

Method	Description
Command line	Use the <code>nbgetconfig</code> command or <code>bpgetconfig</code> command to obtain a list of configuration entries. Then use <code>nbsetconfig</code> or <code>bpsetconfig</code> to change the options as needed. These commands update the appropriate configuration files on both Windows (registry) and UNIX (<code>bp.conf</code> file) primary servers and clients. Use the <code>nbemmcmd</code> command to modify some options on hosts. Detailed information on these commands is available in the NetBackup Commands Reference Guide.
<code>vm.conf</code> file	The <code>vm.conf</code> file contains configuration entries for media and device management. See the NetBackup Administrator's Guide, Volume II for more information.
Backup, Archive, and Restore client interface	Administrators can specify configuration options for NetBackup clients. See the NetBackup Backup, Archive, and Restore Getting Started Guide.

View or edit the host properties of a server or client

The configuration options within the **Host properties** let an administrator customize NetBackup to meet specific site preferences and requirements. The NetBackup web UI displays properties for NetBackup primary servers, media servers, and clients.

Note: In a clustered environment, you must make changes to host properties separately on each node of the cluster.

View or edit the host properties of the primary server

To view or edit the host properties of the primary server

- 1 On the left, click **Hosts > Host properties**.
- 2 At the top left, from the list select **Primary server**.
- 3 Select the primary server and click **Connect**.
- 4 Click **Edit primary server**.
- 5 Make any changes that you want. Then click **Save**.

View or edit the host properties of a media server

To view or edit the host properties of a media server

- 1 On the left, click **Hosts > Host properties**.
- 2 At the top left, from the list select **Media server**.
- 3 Select the media server and click **Connect**.
- 4 Click **Edit media server**.
- 5 Make any changes that you want. Then click **Save**.

View or edit the host properties of a client

To view or edit the host properties of a client

- 1 On the left, click **Hosts > Host properties**.
- 2 At the top left, from the list select **Client server**.
- 3 Select the client and click **Connect**.
- 4 Click **Edit client**.
- 5 Make any changes that you want. Then click **Save**.

Host information and settings in Host properties

In **Hosts > Host properties** you can view the information and certain settings for each host in the NetBackup environment.

Table 7-2 Host properties for hosts

Property name	Description
Host	The NetBackup client name of the host.
Operating system	The operating system and OS version on which the host is installed.
OS type	The type of OS.
Host type	The type of host: Primary server, media server, or client.
IP address	The IP address of the host.
Version	The NetBackup version of the host.

Table 7-2 Host properties for hosts (*continued*)

Property name	Description
Status	Indicates if the host is connected and available for a user to update its host properties. If necessary, select the host and click Connect .
Resiliency	Indicates if Resilient network settings are configured on the primary server. See "Resilient network properties" on page 182.
Host mappings	Lists any host mappings that are configured for the host. See "Approve or add mappings for a host that has multiple host names" on page 434.

Reset a host's attributes

In some cases you need to reset a host's attributes to allow successful communication with the host. A reset is most common when a host is downgraded to a 8.0 or earlier version of NetBackup. After the downgrade, the primary server cannot communicate with the client because the communication status for the client is still set to the secure mode. A reset updates the communication status to reflect the insecure mode.

When you reset a host's attributes:

- NetBackup resets the host ID to host name mapping information, the host's communication status and so on. It does not reset the host ID, host name, or security certificates of the host.
- The connection status is set to the insecure state. The next time the primary server communicates with the host, the connection status is updated appropriately.

To reset the attributes for a host

- 1 On the left, select **Security > Host mappings**.
- 2 Locate the host and click **Actions > Reset attributes**.
- 3 Choose if you want to communicate insecurely with 8.0 and earlier hosts.

NetBackup can communicate with a 8.0 or earlier host when the **Allow communication with NetBackup 8.0 and earlier hosts** option is enabled in the **Global Security Settings**.

Note: If you unintentionally reset a host's attributes, you can undo the changes by restarting the `bpcd` service. Otherwise, the host attributes are automatically updated with the appropriate values after 24 hours.

Active Directory properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the Windows client. If necessary click **Connect**, then and click **Edit client**. Then click **Windows Client > Active Directory**.

The **Active Directory** properties apply to the backup of currently selected Windows Server client. The **Active Directory** properties determine how the backups that allow Active Directory granular restores are performed.

The **Active Directory** host properties contain the following settings.

Table 7-3 Active Directory properties

Property	Description
Perform consistency check before backup when using Microsoft Volume Shadow Copy Service snapshot provider	Checks snapshots for data corruption. Applies only to snapshots that the Microsoft Volume Shadow Copy Services (VSS) performs. If corrupt data is found and this option is not selected, the job fails. See "Windows open file backup tab of the Client attributes properties" on page 110.
Continue with backup if consistency check fails	Continues the backup job even if the consistency check fails. It may be preferable for the job to continue, even if the consistency check fails. For example, a backup of the database in its current state may be better than no backup at all. Or, it may be preferable for the backup of a large database to continue if it encounters only a small problem.

Backup pool host properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Backup host pools**.

The **Backup host pools** properties apply to the backup of the currently selected primary server. A backup host pool is a group of hosts where NetBackup stages the snapshots of the volumes for the backup process to access them. These hosts can be NetBackup clients, media servers, or a primary server.

For the hosts that you add to the backup host pool, their volumes are distributed for backup purposes on the backup hosts. This configuration results in a better backup performance.

You can create a backup host pool with different versions of NetBackup hosts. You can create Windows backup host pools only with version 9.0.1 or later. Windows hosts with a version earlier than 9.0.1 are not displayed.

Note the following important points:

- In a backup host pool you can either have Linux hosts or Windows hosts only. A pool does not support hosts with both platforms.
- All the hosts in the backup host pool must use the same OS version. This way each host has the same version of NFS for consistent backups.
- For backup hosts with a multi-NIC setup, add the host name that is already used on the NetBackup primary server. Do not add an alias name or any other host names in the backup host pool.

Add a backup host pool

To add a backup host pool

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host properties**.
- 3 Select the primary server. If necessary, click **Connect**. Then click **Edit primary server**.
- 4 Click **Backup host pools**.
- 5 Click **Add**.
- 6 Enter the **Backup host pool name**.
- 7 In the **Enter hostname to add to list** box, type the name and click **Add to list**.
- 8 A pool can either have Linux or Windows hosts. To filter the backup hosts in the list, from the **OS type** list select **Windows** or **Linux**.
- 9 From the list, select the hosts that you want to add to the pool.
- 10 Click **Save**.

Add or remove hosts from a backup host pool

To add or remove hosts from a backup host pool

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host properties**.
- 3 Select the primary server. If necessary, click **Connect**. Then click **Edit primary server**.
- 4 Click **Backup host pools**.
- 5 Locate the pool and click **Actions > Edit**.
- 6 A pool can either have Linux or Windows hosts. To filter the backup hosts in the list, from the **OS type** list select **Windows** or **Linux**.
- 7 Select the hosts that you want to include the pool. Or, deselect the hosts you want to remove from the pool.
- 8 Click **Save**.

Delete a backup host pool

You cannot delete a backup host pool if it is part of policy. You must first select a different pool in the policy.

To add or remove hosts from a backup host pool

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host properties**.
- 3 Select the primary server. If necessary, click **Connect**. Then click **Edit primary server**.
- 4 Click **Backup host pools**.
- 5 Locate the pool and click **Actions > Delete > Delete**.

Busy file settings properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the UNIX client. If necessary click **Connect**, then click **Edit client**. Click **UNIX client > Busy file settings**.

The **Busy file settings** properties define what occurs when NetBackup encounters a busy file during a backup of a UNIX client.

The **Busy file settings** host properties contain the following settings.

Table 7-4 Busy file settings properties

Property	Description
Working directory	Specifies the path to the busy-files working directory. On a UNIX client, the value in the user's <code>\$HOME/bp.conf</code> file takes precedence if it exists. By default, NetBackup creates the <code>busy_files</code> directory in the <code>/usr/opensv/netbackup</code> directory.
Administrator email address	Specifies the recipient of the busy-file notification message when the action is set to Send email. By default, the mail recipient is the administrator. On a UNIX client, the value in the user's <code>\$HOME/bp.conf</code> file takes precedence if it exists. By default, <code>BUSY_FILE_NOTIFY_USER</code> is not in any <code>bp.conf</code> file and the mail recipient is <code>root</code> .
Process busy files	Enables busy files to be processed according to the host property settings. NetBackup follows the Busy file settings if it determines that a file changes during a backup. By default, Process busy files is not enabled and NetBackup does not process the busy files. Additional information about busy file processing is available in the NetBackup Administrator's Guide, Volume II .
File action file list	Specifies the absolute path and file name of the busy file. The metacharacters <code>*</code> , <code>?</code> , <code>[]</code> , <code>[-]</code> can be used for pattern matching of file names or parts of file names.
Add	Adds a new file entry. Enter the file and path directly, or browse to select a file.
Actions > Delete	Deletes the selected file from the file action list.
Retry count	Specifies the number of times to try the backup. The default retry count is 1.
Busy file action	The following options specify which action to take when busy-file processing is enabled. On a UNIX client, the value in the user's <code>\$HOME/bp.conf</code> file takes precedence if it exists. ■ Send email sends a busy sends a busy file notification message to the user that is specified in Administrator email address. ■ Retry the backup retries the backup on the specified busy file. The Retry count value determines the number of times NetBackup tries a backup. ■ Ignore excludes the busy file from busy file processing. The file is backed up, then a log entry that indicates it was busy appears in the All log entries report. The All logs entries report is also available in the NetBackup web UI.

Activating the Busy file settings in host properties

To activate the settings in the **Busy file settings** host properties, use the following procedure.

To activate Busy file settings

- 1 Copy the `bpend_notify_busy` script:

```
/usr/opensv/netbackup/bin/goodies/bpend_notify_busy
```

to the path:

```
/usr/opensv/netbackup/bin/bpend_notify
```

- 2 Set the file access permissions to allow group and others to run `bpend_notify`.
- 3 Configure a policy with a user backup schedule for the busy file backups.

This policy services the backup requests that the repeat option in the actions file generates. The policy name is significant. By default, NetBackup alphabetically searches (uppercase characters first) for the first available policy with a user backup schedule and an open backup window. For example, a policy name of `AAA_busy_files` is selected ahead of `B_policy`.

Clean up properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Clean up**.

The **Clean up** properties manage the retention of various logs and incomplete jobs. The **Clean up** properties apply to primary servers.

The **Clean up** host properties contain the following settings.

Table 7-5 Clean up properties

Property	Description
Keep true image restoration (TIR) information	Specifies the number of days to keep true image restore information on disk. After the specified number of days, the images are pruned (removed). Applies to all policies for which NetBackup collects true image restore information. The default is one day. When NetBackup performs a true image backup, it stores the following images on the backup media: ■ Backed up files ■ True image restore information NetBackup also stores the true image restore information on disk in the following directories: On Windows: <pre>install_path\NetBackup\db\images</pre> On UNIX: <pre>/usr/opensv/netbackup/db/images</pre> NetBackup retains the information for the number of days that this property specifies. Keeping the information on disk speeds up restores. If a user requests a true image restore after the information was deleted from disk, NetBackup retrieves the required information from the media. The only noticeable difference to the user is a slight increase in total restore time. NetBackup deletes the additional information from disk again after one day.
Move restore job from incomplete state to done state	Indicates the number of days that a failed restore job can remain in an Incomplete state. After that time, the Activity monitor shows the job as Done. The default is 7 days. The maximum setting is 365 days. If Checkpoint Restart for restores is used, the Restore retries property allows a failed restore job to be retried automatically. See "Universal settings properties" on page 212.

Table 7-5 Clean up properties (*continued*)

Property	Description
Move backup job from incomplete state to done state	Indicates the maximum number of hours that a failed backup job can remain in an incomplete state. After that time, the Activity Monitor shows the job as Done. The minimum setting is 1 hour. The maximum setting is 72 hours. The default is 3 hours. When an active job has an error, the job goes into an Incomplete state. In the Incomplete state, the administrator can correct the condition that caused the error. If an Incomplete job does not complete successfully and is moved to the Done state, the job retains the error status. Note: A resumed job reuses the same job ID, but a restarted job receives a new job ID. The job details indicate that the job was resumed or restarted. Note: This property does not apply to suspended jobs. Suspended jobs must be resumed manually before the retention period of the job is met and the image expires. If a suspended job is resumed after the retention period is met, the job fails and is moved to the Done state.
Image cleanup interval	Specifies the maximum interval that can elapse before an image cleanup is run. Image cleanup is run after every successful backup session (that is, a session in which at least one backup runs successfully). If a backup session exceeds this maximum interval, an image cleanup is initiated.
Catalog cleanup wait time	Specifies the minimum interval that can elapse before an image cleanup is run. Image cleanup is not run after a successful backup session until this minimum interval has elapsed since the previous image cleanup.

Client name properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the client. If necessary click **Connect**, then click **Edit client**. Click **Client name**.

The **Client name** property specifies the NetBackup client name for the selected client. The name must match the name the policy uses to back up the client. The only exception is for a redirected restore, where the name must match that of the client whose files are to be restored. The client name is initially set during installation.

The name that is entered here must also match the client name in the **Client attributes** for the primary server. If it does not match, the client cannot browse for its own backups.

Note: Using an IPv6 address as a client name in a policy can cause backups to fail. Specify a host name instead of an IPv6 address.

See [“Client attributes properties”](#) on page 104.

If the value is not specified, NetBackup uses the name that is set in the following locations:

- For a Windows client
In the Network application from the Control Panel.
- For a UNIX client
The name that is set by using the `hostname` command.
The name can also be added to a `$HOME/bp.conf` file on a UNIX client. However, the name is normally added in this manner only for redirected restores. The value in the `$HOME/bp.conf` file takes precedence if it exists.

Client attributes properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Client attributes**.

The **Client attributes** properties apply to the clients of currently selected primary server.

The **Global client attributes** property applies to all clients, unless overridden as described in the following table.

Table 7-6 Global client attributes

Attribute	Description
Allow client browse	Allows all clients to browse files for restoring. This attribute is overridden if the Browse and restore ability option on the General tab is set to Deny both for a particular clients.
Allow client restore	Allows all clients to restore files. This attribute is overridden if the Browse and restore ability option on the General tab is set to Allow browse only or Deny both .

Table 7-6 Global client attributes (*continued*)

Attribute	Description
Clients	Specifies the list of clients in the client database on the currently selected primary server. A client must be in the client database before you can change the client properties in Client attributes. The client database consists of directories and files in the following directories: Windows: <code>install_path\NetBackup\db\client</code> UNIX: <code>/usr/opensv/netbackup/db/client</code> If a client is not listed in the Clients list, click Add to add a client to the client database. Enter a client name in the text box or select a client. Then click Add. The name that is entered here must match the Client name property for the specific client. If it does not match, the client cannot browse its own backups. See “Client name properties” on page 103. Use the <code>bpclient</code> command to add clients to the client database if dynamic addressing (DHCP) is in use. Additional information about busy file processing is available in the NetBackup Administrator's Guide, Volume II. On UNIX: You also can create, update, list, and delete client entries by using the <code>bpclient</code> command that is located in the following directory: <code>/usr/opensv/netbackup/bin/admincmd</code>
General tab	Specifies how to configure the selected Windows primary servers (clients). See “General tab of the Client attributes properties” on page 106.
Connect options tab	Specifies how to configure the connection between a NetBackup server and a NetBackup client. See “Connect options tab of the Client attributes properties” on page 109.
Windows open file backup tab	Specifies whether a client uses Windows Open File Backup. Also, specifies whether Volume Snapshot Provider or Volume Shadow Copy Service is used as the snapshot provider. See “Windows open file backup tab of the Client attributes properties” on page 110.

General tab of the Client attributes properties

To access this tab, in the web UI select **Hosts > Host properties**. Select the Windows primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Client attributes**. Then click the **General** tab.

The properties on the **General** tab apply to selected Windows primary servers. The tab appears on the **Client attributes** page.

The **General** tab contains the following properties.

Table 7-7 General tab properties

Property	Description
Disable backups until:	Makes the specified clients in the General tab unavailable for backups until the specified date and time. By default, clients are online and included in the policies in which they are listed. When Disable backups until is selected for a client, no jobs are scheduled for that client. Since the client is not part of any job, no backup status is listed for the client. If a client is taken offline, any job is allowed to complete that includes the client and is already running. If a backup or restore job is manually submitted for a client that is offline, the Activity monitor displays the job as failed with a status code 1000 (Client is offline). Note: Changes to this property do not appear in the audit report. The ability to take clients offline is useful in a number of situations. See “Offline option usage considerations and restrictions” on page 108.
Disable restores until:	Makes the specified clients in the General tab unavailable for restores until the specified date and time. By default, clients are online and available for restore.

Table 7-7 **General tab properties** (*continued*)

Property	Description
Maximum data streams	Specifies the maximum number of jobs that are allowed at one time for each selected client. (This value applies to the number of jobs on the client, even if multistreaming is not used.) To change the setting, select Maximum data streams. Then scroll to or enter a value up to 99. The Maximum data streams property interacts with Maximum jobs per client and Limit jobs per policy as follows: ■ If the Maximum data streams property is not set, the limit is either the one indicated by the Maximum jobs per client property or the Limit jobs per policy property, whichever is lower. ■ If the Maximum data streams property is set, NetBackup uses either Maximum jobs per client property. NetBackup uses either Maximum data streams or Limit jobs per policy, whichever is lower. See “Global attributes properties” on page 147.
Browse and restore	Specifies the client permissions to list and restore backups and archives. Select the clients in the General tab of the Client attributes and choose a Browse and restore property. To use the Global client attributes settings, select Use global settings. ■ To allow users on the selected clients to both browse and restore, select Allow both. ■ To allow users on the selected clients to browse but not restore, select Allow browse only. ■ To prevent users on the selected clients from the ability to browse or restore, select Deny both.
Browse and restore scheduled backups	Specifies whether the clients can list and restore from scheduled backups. (This setting does not affect user backups and archives.) This property applies to the privileges that are allowed to a non-Windows administrator or non-root user who is logged into the client. This property also applies to the users that do not have backup and restore privileges. Windows administrators and root users can list and restore from scheduled backups as well as user backups regardless of the Browse and restore scheduled backups setting.
Deduplication	Specifies the deduplication action for clients if you use the NetBackup Data Protection Optimization Option. For a description of the client-side deduplication options and their actions: See “Where deduplication should occur” on page 108.

Offline option usage considerations and restrictions

The ability to take clients offline is useful in a number of situations. For example, in the event of planned outages or maintenance, client systems can be taken offline to avoid the unnecessary errors that administrators would then need to investigate. This option can also be used to anticipate new clients in the system. You can add them to policies but configure them as offline until they are in place and ready to use.

The following actions can be performed if a client is offline.

Table 7-8 Offline option actions

Type of job or operation	Action or restriction
A client is offline and the job is already in progress.	Offline clients continue to be included in any job.
A client is offline and job retries were started before the client was taken offline.	Job retries continue as normal.
Any duplication job that is associated with a storage lifecycle policy and an offline client.	Continues to run until complete.
Restore jobs	Can be run for offline clients.
The user attempts a manual backup for an offline client.	The backup fails with a status code 1000 (Client is offline). The user can either wait until the client is brought online again or bring the client online manually. Use either the NetBackup web UI or the <code>bpcclient</code> command to do so before resubmitting the manual job.
Archive backups	Not allowed for offline clients.
Administrators restarting or resuming jobs.	Not allowed for offline clients.

Caution: If the primary server is offline, hot catalog backups cannot run.

Where deduplication should occur

The **Deduplication** property specifies the deduplication action for clients if you use the NetBackup Data Protection Optimization Option. More information is available on the client-side deduplication options.

See [Table 7-9](#) on page 109.

The primary server and the clients (that deduplicate their own data) must use the same name to resolve the storage server. The name must be the host name under which the NetBackup Deduplication Engine credentials were created. If they do not

use the same name, backups fail. In some environments, careful configuration may be required to ensure that the client and the primary server use the same name for the storage server. Such environments include those that use VLAN tagging and those that use multi-homed hosts.

NetBackup does not support the following for client-side deduplication:

- Multiple copies per each job configured in a NetBackup backup policy. For the jobs that specify multiple copies, the backup images are sent to the storage server and may be deduplicated there.
- NDMP hosts. The backup jobs fail if you try to use client-side deduplication for NDMP hosts.

Table 7-9 Client-side deduplication options

Option	Description
Always use the media server (the default)	Always deduplicates the data on the media server. The default. Jobs fail if one of the following is true: ■ The deduplication services on the storage server are inactive. ■ The deduplication pool is down.
Prefer to use client-side deduplication	Deduplicates the data on the client and then sends it directly to the storage server. NetBackup first determines if the storage server is active. If it is active, the client deduplicates the backup data and sends it to the storage server to be written to disk. If it is not active, the client sends the backup data to a media server, which deduplicates the data.
Always use client-side deduplication	Always deduplicates the backup data on the client and then sends it directly to the storage server. If a job fails, NetBackup does not retry the job.

You can override the **Prefer to use client-side deduplication** or **Always use client-side deduplication** host property in the backup policies.

More information about client deduplication is available in the [NetBackup Deduplication Guide](#).

Connect options tab of the Client attributes properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server. If necessary click **Connect**, then click **Edit primary server**. Click **Client attributes**. Then click the **Connect options** tab.

The properties in the **Connect options** tab describe how a NetBackup server connects to NetBackup clients. The tab appears on the **Client attributes** page.

The **Connect options** tab contains the following options.

Table 7-10 Connect options tab properties

Property	Description
BPCD connect back	Specifies how daemons are to connect back to the NetBackup Client daemon (BPCD) and contains the following options: ■ Use default connect options Uses the value that is defined in the Firewall host properties of the client's NetBackup server. See "Firewall properties" on page 142. ■ Random port NetBackup randomly chooses a free port in the allowed range to perform the legacy connect-back method. ■ VNETD port NetBackup uses the <code>vnetd</code> port number for the connect-back method.
Ports	Specifies the method that the selected clients should use to connect to the server and contains the following options: ■ Use default connect options Uses the value that is defined in the Firewall host properties of the client's NetBackup server. See "Firewall properties" on page 142. ■ Reserved ports Uses a reserved port number. ■ Non-reserved ports Uses a non-reserved port number.

Windows open file backup tab of the Client attributes properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the Windows primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Client attributes**. Then click the **Windows open file backup** tab.

Use the settings in this tab only if you want to change the default settings.

By default, NetBackup uses Windows open file backups for all Windows clients. (No clients are listed in the **Client attributes** page.) The server uses the following default settings for all Windows clients:

- Windows open file backup is enabled on the client.

- Microsoft Volume Shadow Copy Service (VSS).
- Snapshots are taken of individual drives (**Individual drive snapshot**) as opposed to all drives at once (**Global drive snapshot**).
- Upon error, the snapshot is terminated (**Abort backup on error**).

Snapshots are a point-in-time view of a source volume. NetBackup uses snapshots to access busy or active files during a backup job. Without a snapshot provider, active files are not accessible for backup.

Table 7-11 Windows open file backup tab properties

Property	Description
Add	Adds a NetBackup client to the list, if you want to change the default settings for Windows open file backups.
Delete	Deletes a client from the list.
Enable Windows open file backup for the selected client	Specifies that Windows open file backup is used for the selected clients. This option functions independently from the Perform Snapshot backups policy option that is available when the Snapshot Client is licensed. If a client is included in a policy that has the Perform Snapshot backups policy option disabled and you do not want snapshots, the Enable Windows open file backups for this client property must be disabled as well for the client. If both options are not disabled, a snapshot is created, though that may not be the intention of the administrator.
Snapshot Provider	Selects the snapshot provider for the selected clients: ■ Use Veritas Volume Snapshot Provider (VSP) This option is used for back-level versions of NetBackup only. Support for those client versions has ended. ■ Use Microsoft Volume Shadow Copy Service (VSS) Uses VSS to create volume snapshots of volumes and logical drives for the selected clients. For information about how to do Active Directory granular restores when using VSS, see the following topic: See “Active Directory properties” on page 97.

Table 7-11 Windows open file backup tab properties (*continued*)

Property	Description
Snapshot usage	Note: The Individual drive snapshot property and the Global drive snapshot property only apply to the non-multistreamed backups that use Windows open file backup. All multistreamed backup jobs share the same volumes snapshots for the volumes in the multistreamed policy. The volume snapshots are taken in a global fashion. Selects how snapshots are made for the selected clients: Individual drive snapshot Specifies that the snapshot should be of an individual drive (default). When this property is enabled, snapshot creation and file backup are done sequentially on a per volume basis. For example, assume that drives C and D are backed up. If the Individual drive snapshot property is selected, NetBackup takes a snapshot of drive C, backs it up, and discards the snapshot. NetBackup then takes a snapshot of drive D, backs it up, and discards the snapshot. Volume snapshots are enabled on only one drive at a time, depending on which drive is to be backed up. This mode is useful when relationships do not have to be maintained between files on the different drives. Global drive snapshot Specifies that the snapshot is of a global drive. All the volumes that require snapshots for the backup job (or stream group for multistreamed backups) are taken at one time. If snapshot creation is not successful, use the Individual drive snapshot option. For example, assume that drives C and D are to be backed up. In this situation, NetBackup takes a snapshot of C and D. Then NetBackup backs up C and backs up D. NetBackup then discards the C and D snapshots. This property maintains file consistency between files in different volumes. The backup uses the same snapshot that is taken at a point in time for all volumes in the backup.

Table 7-11 Windows open file backup tab properties (*continued*)

Property	Description
Snapshot error control	Determines the action to take if there is a snapshot error: Abort backup on error Stops the backup if there is an error during the backup job (after the snapshot is created). The most common reason for a problem after the snapshot is created and is in use by a backup, is that the cache storage is full. If the Abort backup on error property is selected (default), the backup job cancels with a snapshot error status if the backup detects a snapshot issue. This property does not apply to successful snapshot creation. The backup job continues regardless of whether a snapshot was successfully created for the backup job. Disable snapshot and continue Destroys the volume snapshots if the snapshot becomes invalid during a backup. The backup continues with Windows open file backups disabled. Regarding the file that had a problem during a backup—it may be that the file was not backed up by the backup job. The file may not be able to be restored. Note: Volume snapshots typically become invalid during the course of a backup because insufficient cache storage was allocated for the volume snapshot. Reconfigure the cache storage configuration of the Windows open file backup snapshot provider to a configuration that best suits your client's installation.

Client settings properties for UNIX clients

To access this setting, in the web UI select **Hosts > Host properties**. Select the UNIX client. If necessary click **Connect**, then click **Edit client**. Click **UNIX client > Client settings**.

The UNIX **Client settings** properties apply to currently selected NetBackup client running on the UNIX platform.

The UNIX **Client settings** host properties contain the following settings.

Table 7-12 UNIX Client settings properties

Property	Description
Locked file action	Determines what happens when NetBackup tries to back up a file with mandatory file locking enabled in its file mode. Select one of the following options: ■ Wait By default, NetBackup waits for files to become unlocked. If the wait exceeds the Client read timeout host property that is configured on the primary server, the backup fails with a status 41. See "Timeouts properties" on page 209. ■ Skip NetBackup skips the files that currently have mandatory locking set by another process. A message is logged if it was necessary to skip a file.
File compression memory	Specifies the amount of memory available on the client when files are compressed during backup. If you select compression, the client software uses this value to determine how much space to request for the compression tables. The more memory that is available to compress code, the greater the compression and the greater the percentage of computer resources that are used. If other processes also need memory, use a maximum value of half the actual physical memory on a computer to avoid excessive swapping. The default is 0. This default is reasonable; change it only if problems are encountered.
Reset file access time to the value before backup	Specifies that the access time (<i>atime</i>) for a file displays the backup time. By default, NetBackup preserves the access time by resetting it to the value it had before the backup. Note: This setting affects the software and the administration scripts that examine a file's access time. Note: If NetBackup Accelerator is used to perform the backup, this setting is ignored. Accelerator does not record and reset the <i>atime</i> for the files that it backs up.
Keep status of user-directed backups, archives, and restores	Specifies the number of days to keep progress reports before the reports are deleted. The default is 3 days. The minimum is 0. The maximum is 9,999 days. Logs for user-directed operations are stored on the client system in the following directory: <pre>install_path\NetBackup\logs\user_ops\loginID\logs</pre>

Table 7-12 UNIX Client settings properties (*continued*)

Property	Description
Use VxFS File Change Log (FCL) for incremental backups	Determines if NetBackup uses the File Change Log on VxFS clients. The default is off. See “VxFS file change log (FCL) for incremental backups property” on page 115.
Default cache device path for snapshots	This setting identifies a raw partition available to the copy-on-write process. This raw partition is used when either nbu_snap or VxFS_Snapshot are selected as the snapshot method. The partition must exist on all the clients that are included in the policy.
Add	Adds the file endings to the list of file endings that you do not want to compress. Click Add, then type the file extension. Click Add to add the ending to the list.
Do not compress files ending with these file extensions	Specifies a list of file extensions. During a backup, NetBackup does not compress files with these extensions because the file may already be in a compressed format. Do not use wildcards to specify these extensions. For example, <code>.A1</code> is allowed, but not <code>.A*</code> or <code>.A[1-9]</code> Files that are already compressed become slightly larger if compressed again. If compressed files with a unique file extension already exist on a UNIX client, exclude it from compression by adding it to this list. Corresponds to adding a <code>COMPRESS_SUFFIX =.suffix</code> option to the <code>bp.conf</code> file.

VxFS file change log (FCL) for incremental backups property

The **Use VxFS File Change Log (FCL) for incremental backups** property is supported on all platforms and versions where VxFS file systems support FCL.

The following VxFS file systems support FCL:

- Solaris SPARC platform running VxFS 4.1 or later.
- AIX running VxFS 5.0 or later.
- HP 11.23 running VxFS 5.0 or later.
- Linux running VxFS 4.1 or later.

The File Change Log (FCL) tracks changes to files and directories in a file system. Changes can include the following: files created; links and unlinks; files renamed; data that is appended; data that is overwritten; data that is truncated; extended attribute modifications; holes punched; and file property updates.

NetBackup can use the FCL to determine which files to select for incremental backups, which can potentially save unnecessary file system processing time. The FCL information that is stored on each client includes the backup type, the FCL offset, and the timestamp for each backup.

The advantages of this property depend largely on the number of file system changes relative to the file system size. The performance effect of incremental backups ranges from many times faster or slower, depending on file system size and use patterns.

For example, enable this property for a client on a very large file system that experiences relatively few changes. The incremental backups may complete sooner for the client since the policy needs to read only the FCL to determine what needs to be backed up on the client.

If a file experiences many changes or multiple changes to many files, the time saving benefit may not be as great.

The following items must be in place for the **Use VxFS File Change Log (FCL) for incremental backups** property to work:

- Enable the **Use VxFS File Change Log (FCL) for incremental backups** property for every client that wants NetBackup to take advantage of the FCL.
- Enable the FCL on the VxFS client.
See the Veritas Storage Foundation documentation for information about how to enable the FCL on the VxFS client.
- Enable the **Use VxFS File Change Log (FCL) for incremental backups** property on the clients in time for the first full backup. Subsequent incremental backups need this full backup to stay synchronized.
- Specify the VxFS mount point in the policy backup selections list in one of the following ways:
 - Specify ALL_LOCAL_DRIVES.
 - Specifying the actual VxFS mount point.
 - Specifying a directory at a higher level than the VxFS mount point, provided that option **Cross mount points** is enabled.

If the policy has **Collect true image restore information** or **Collect true image restore information with move detection** enabled, it ignores the **Use VxFS File Change Log (FCL) for incremental backups** property on the client.

The following table describes the additional options that are available on the VxFS file change log feature.

Table 7-13 VxFS file change log feature options

Option	Description
Activity monitor messages	Displays any messages that note when the file change log is used during a backup as follows: Using VxFS File Change Log for backup of <i>pathname</i> Also notes when full and incremental backups are not synchronized.
Keeping the data files synchronized with the FCL	The data files must be in sync with the FCL for this property to work. To keep the data files synchronized with the FCL on the VxFS client do not turn off the FCL and turn it on again. Note: If NetBackup encounters any errors as it processes the FCL, it switches to the normal files system scan. If this switch occurs, it appears in the Activity monitor.
VxFS administration	Additional VxFS commands are available to administrate the FCL in the Veritas Storage Foundation documentation.

Client settings properties for Windows clients

To access these settings, in the web UI select **Hosts > Host properties**. Select the Windows client and click **Edit client**. Then click **Windows client > Client settings**.

The Windows **Client settings** properties apply to the currently selected Windows client .

The **Windows clients > Client settings** host properties contain the following settings.

Table 7-14 Client settings properties for Windows clients

Property	Description
General level	Enables logs for <code>bpnetd</code> , <code>bpbkar</code> , <code>tar</code> , and <code>nbwin</code> . The higher the level, the more information is written. The default is Minimum logging .

Table 7-14 Client settings properties for Windows clients (*continued*)

Property	Description
Wait time before clearing archive bit	Specifies how long the client waits before the archive bits for a differential incremental backup are cleared. The minimum allowable value is 300 (default). The client waits for acknowledgment from the server that the backup was successful. If the server does not reply within this time period, the archive bits are not cleared. This option applies only to differential-incremental backups. Cumulative-incremental backups do not clear the archive bit.
Use Windows change journal	Note: The Use Windows Change Journal option applies to Windows clients only. This option works together with the Use Accelerator policy attribute and the Accelerator forced rescan schedule attribute.
Time overlap	Specifies the number of minutes to add to the date range for incremental backups when you use date-based backups. This value compensates for differences in the speed of the clock between the NetBackup client and server. The default is 60 minutes. This value is used during incremental backups when you use the archive bit and when you examine the create time on folders. This comparison is done for archive bit-based backups as well as date-based backups.
Communications buffer size	Specifies the size (in kilobytes) of the TCP and the IP buffers that NetBackup uses to transfer data between the NetBackup server and client. For example, specify 10 for a buffer size of 10 kilobytes. The minimum allowable value is 2, with no maximum allowable value. The default is 128 kilobytes.
User-directed timeouts	Specifies the seconds that are allowed between when a user requests a backup or restore and when the operation begins. The operation fails if it does not begin within this time period. This property has no minimum value or maximum value. The default is 60 seconds.
Perform default search for restore	Instructs NetBackup to search the default range of backup images automatically. The backed up folders and files within the range appear whenever a restore window is opened. Clear the Perform default search for restore check box to disable the initial search. With the property disabled, the NetBackup Restore window does not display any files or folders upon opening. The default is that the option is enabled.

Table 7-14 Client settings properties for Windows clients (*continued*)

Property	Description
TCP level	Enables logs for TCP. Scroll to one of the following available log levels: ■ 0 No extra logging (default) ■ 1 Log basic TCP/IP functions ■ 2 Log all TCP/IP functions ■ 3 Log contents of each read/write Note: Setting the TCP level to 2 or 3 can cause the status reports to be very large. It can also slow a backup or restore operation.
Incrementals	■ Based on timestamp Files that are selected for backup based on the date that the file was last modified. When Use change journal is selected, Based on timestamp is automatically selected. ■ Based on archive bit Note: It is not recommended that you combine differential incremental backups and cumulative incremental backups within the same Windows policy when the incremental backups are based on archive bit. NetBackup include files in an incremental backup only if the archive bit of the file is set. The system sets this bit whenever a file is changed and it normally remains set until NetBackup clears it. A full backup always clears the archive bit. A differential-incremental backup clears the archive bit if the file is successfully backed up. The differential-incremental backup must occur within the number of seconds that the Wait time before clearing archive bit property indicates. A cumulative-incremental or user backup has no effect on the archive bit. If you install or copy files from another computer, the new files retain the date timestamp of the originals. If the original date is before the last backup date on this computer, then the new files are not backed up until the next full backup.
Maximum error messages for single issue	Defines how many times a NetBackup client can send the same error message to a NetBackup server. For example, if the archive bits cannot be reset on a file, this property limits how many times the message appears in the server logs. The default is 10.
Keep status of user-directed backups, archives and restores	Specifies how many days the system keeps progress reports before NetBackup automatically deletes them. The default is 3 days.

How to determine if change journal support is useful in your NetBackup environment

Using NetBackup support for the change journal is beneficial only where the volumes are large and relatively static.

Suitable candidates for enabling NetBackup change journal support are as follows:

- If the NTFS volume contains more than 1,000,000 files and folders and the number of changed objects between incremental backups is small (less than 100,000), the volume is a good candidate for enabling NetBackup change journal support.

Unsuitable candidates for enabling NetBackup change journal support are as follows:

- Support for the change journal is intended to reduce scan times for incremental backups by using the information that is gathered from the change journal on a volume. Therefore, to enable NetBackup change journal support is not recommended if the file system on the volume contains relatively few files and folders. (For example, hundreds of thousands of files and folders.) The normal file system scan is suitable under such conditions.
- If the total number of changes on a volume exceeds from 10% to 20% of the total objects, the volume is not a good candidate for enabling NetBackup change journal support.
- Be aware that virus scanning software can interfere with the use of the change journal. Some real-time virus scanners intercept a file open for read, scan for viruses, then reset the access time. This results in the creation of a change journal entry for every scanned file.

Guidelines for enabling NetBackup change journal support

The following items are guidelines to consider for enabling NetBackup change journal support:

- Change journal support is not offered for user-directed backups. The USN stamps for full and incremental backups in the permanent record do not change.
- NetBackup support for change journal works with checkpoint restart for restores.
- Support for change journal is not offered with several NetBackup options.

If **Use Windows change journal** is enabled, it has no effect while you use the following options or products:

- True image restore (TIR) or True image restore with Move Detection
- Synthetic backups
- Bare Metal Restore (BMR)

For more information, see the *NetBackup Bare Metal Restore Administrator's Guide*.

See [“How to determine if change journal support is useful in your NetBackup environment”](#) on page 120.

Cloud Storage properties

Note: To access these properties, in the web UI select **Hosts > Host properties**. Select the primary server and click **Edit primary server**. Then click **Cloud Storage**.

The NetBackup **Cloud Storage** properties apply to the currently selected primary server.

The hosts that appear in this **Cloud Storage** list are available to select when you configure a storage server. The **Service provider** type of your cloud vendor determines whether a service host is available or required.

NetBackup includes service hosts for some cloud storage providers. You can add a new host to the **Cloud Storage** list if the **Service provider** type allows it. If you add a host, you also can change its properties or delete it from the **Cloud Storage** list. (You cannot change or delete the information that is included with NetBackup.)

If you do not add a service host to this **Cloud Storage** list, you can add one when you configure the storage server. The **Service provider** type of your cloud vendor determines whether a **Service host name** is available or required.

Cloud Storage host properties contain the following properties:

Table 7-15 Cloud Storage

Property	Description
Cloud Storage	The cloud storage that corresponds to the various cloud service providers that NetBackup supports are listed here.
Associated cloud storage servers for <host>	The cloud storage servers that correspond to the selected cloud storage are displayed.

For more information about NetBackup cloud storage, see the [NetBackup Cloud Administrator's Guide](#).

Credential access properties

Note: To access these settings, in the web UI select **Hosts > Host properties**. Select the primary server and click **Edit primary server**. Then click **Credential access**.

Certain NetBackup hosts that are not named as clients in a policy must be enabled to access NDMP or disk array credentials. Use the **Credential access** properties to enter the names of those NetBackup hosts.

The **Credential access** host properties contain the following settings.

Table 7-16 Credential access host properties

Property	Description
NDMP Clients list	To add an NDMP client to the NDMP clients list, click Add . Enter the names of the NDMP hosts that are not named as clients in a policy.
Disk clients list	To add a disk client to the Disk clients list, click Add. Enter the names of the NetBackup hosts that meet all of the following criteria: ■ The host must be designated in a policy as the Off-host backup host in an alternate client backup. ■ The host that is designated as the off-host backup computer must not be named as a client on the Clients tab in any NetBackup policy. ■ The policy for the off-host backup must be configured to use one of the disk array snapshot methods for the EMC CLARiiON, HP EVA, or IBM disk arrays. Note: The credentials for the disk array or NDMP host are specified in the NetBackup web UI. Click Credential management and then click on the Client credentials tab. Note: Off-host alternate client backup is a feature of NetBackup Snapshot Client, which requires a separate license. The NetBackup for NDMP feature requires the NetBackup for NDMP license.

Data Classification properties

To access these settings, in the web UI click **Hosts > Host properties**. Select the server and click **Edit media server** or **Edit primary server**. Then click **Data classification**.

The **Data classification** properties apply to currently selected primary or media server.

Data classifications must be configured in the **Data classification** host properties before storage lifecycle policies can be configured.

Note: Data classifications cannot be deleted. However, the name, description, and the rank can be changed. The classification ID remains the same.

The **Data classification** page contains the following properties.

Table 7-17 Data classification properties

Property	Description
Rank column	The Rank column displays the rank of the data classifications. The order of the data classifications determines the rank of the classification in relationship to the others in the list. The lowest numbered rank has the highest priority. Use the Up and Down buttons to move the classification up or down in the list. To create a new data classification, click Add. New data classifications are added to bottom of the list.
Name column	The Name column displays the data classification name. While data classifications cannot be deleted, the data classification names can be modified. NetBackup provides the following data classifications by default: ■ Platinum (highest rank by default) ■ Gold (second highest rank by default) ■ Silver (third highest rank by default) ■ Bronze (lowest rank by default)
Description column	In the Description, enter a meaningful description for the data classification. Descriptions can be modified.
Data Classification ID	The Data classification ID is the GUID value that identifies the data classification and is generated when a new data classification is added and the host property is saved. . A data classification ID becomes associated with a backup image by setting the Data classification attribute in the policy. The ID is written into the image header. The storage lifecycle policies use the ID to identify the images that are associated with classification. ID values can exist in image headers indefinitely, so data classifications cannot be deleted. The name, description, and rank can change without changing the identity of the data classification.

Add a data classification

Use the following procedures to create or change a data classification.

To add a data classification

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host properties**.
- 3 Click **Data classification**.
- 4 Click **Add**.
- 5 Add the name and description.
- 6 Click **Add**.

Note: Data classifications cannot be deleted.

- 7 To change the priority of a classification, select a row and click **Up** or **Down** options.

Default job priorities properties

To access these settings, in the web UI select **Hosts > Host properties**. Select the primary server and click **Edit primary server**. Then click **Default job priorities**.

The **Default job priorities** host properties let administrators configure the default job priority for different job types.

The job priority can be set for individual jobs in the following utilities:

- In the **Jobs** tab of the **Activity monitor** for queued or active jobs.
- In the **Catalog** utility for verify, duplicate, and import jobs.
- In the **Backup, Archive, and Restore** client interface for restore jobs.

The **Default job priorities** page contains the following properties.

Table 7-18 Default job priorities properties

Property	Description
Job type	The type of job.

Table 7-18 Default job priorities properties (*continued*)

Property	Description
Job priority	The priority that a job has as it competes with other jobs for backup resources. The value can range from 0 to 99999. The higher the number, the greater the priority of the job. A new priority setting affects all the policies that are created after the host property has been changed. A higher priority does not guarantee that a job receives resources before a job with a lower priority. NetBackup evaluates jobs with a higher priority before those with a lower priority. However, the following factors can cause a job with a lower priority to run before a job with a higher priority: ■ To maximize drive use, a low priority job may run first if it can use a drive that is currently loaded. A job with a higher priority that requires that the drive be unloaded would wait. ■ If a low priority job can join a multiplexed group, it may run first. The job with a higher priority may wait if it is not able to join the multiplexed group. ■ If the NetBackup Resource Broker (<code>nbrb</code>) receives a job request during an evaluation cycle, it does not consider the job until the next cycle, regardless of the job priority.

Understanding the job priority setting

NetBackup uses the **Job priority** setting as a guide. Requests with a higher priority do not always receive resources before a request with a lower priority.

NetBackup evaluates the requests sequentially and sorts them based on the following criteria:

- The request's first priority.
- The request's second priority.
- The birth time (when the Resource Broker receives the request).

The first priority is weighted more heavily than the second priority, and the second priority is weighted more heavily than the birth time.

Because a request with a higher priority is listed in the queue before a request with a lower priority, the request with a higher priority is evaluated first. Even though the chances are greater that the higher priority request receives resources first, it is not always definite.

The following scenarios present situations in which a request with a lower priority may receive resources before a request with a higher priority:

- A higher priority job needs to unload the media in a drive because the retention level (or the media pool) of the loaded media is not what the job requires. A lower priority job can use the media that is already loaded in the drive. To maximize drive utilization, the Resource Broker gives the loaded media and drive pair to the job with the lower priority.
- A higher priority job is not eligible to join an existing multiplexing group but a lower priority job is eligible to join the multiplexing group. To continue spinning the drive at the maximum rate, the lower priority job joins the multiplexing group and runs.
- The Resource Broker receives resource requests for jobs and places the requests in a queue before it processes them. New resource requests are sorted and evaluated every 5 minutes. Some external events (a new resource request or a resource release, for example) can also start an evaluation. If the Resource Broker receives a request of any priority while it processes requests in an evaluation cycle, the request is not evaluated until the next evaluation cycle starts.

Distributed application restore mapping properties

To access these settings, in the web UI click **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Distributed application restore mapping**.

Some applications, such as SharePoint, Exchange, and SQL Server distribute and replicate data across multiple hosts. Or, the configuration includes a cluster where communication occurs across multiple nodes. Use the **Distributed application restore mapping** to provide a mapping of the hosts in the database environment so that NetBackup can successfully restore the databases. See the administrator's guide for the database agent for more details.

For example, for a SharePoint farm assume that the farm has two application servers (App1 and App2), one front-end server (FE1) and one SQL database (SQLDB1). The Distributed application restore mapping for this SharePoint server would be as following follows:

Application host	Component host
App1	SQLDB1
App2	SQLDB1
FE1	SQLDB1

The **Distributed application restore mapping** page contains the following properties.

Table 7-19 Distributed application restore mapping properties

Property	Description
Add	This option adds a component host that is authorized to run restores on a SharePoint, Exchange, or SQL Server application host. For SharePoint, NetBackup catalogs backup images under the front-end server name. To allow NetBackup to restore SQL Server back-end databases to the correct hosts in a farm, provide a list of the SharePoint hosts. For Exchange, any operations that use Granular Recovery Technology (GRT) require that you provide a list of the Exchange virtual and the physical host names. You must also include the off-host client and the granular proxy host. For SQL Server, this configuration is required for restores of a SQL Server cluster or a SQL Server availability group (AG). Note: For VMware backups and restores that protect SharePoint, Exchange, or SQL Server, you only need to add the hosts that browse for backups or perform restores. You must also configure a mapping if you use a Primary VM Identifier other than the VM hostname. See the administrator's guide for the database agent for more details. Note: Use either the client's short name or its fully qualified domain name (FQDN). You do not need to provide both names in the list. For more details, see the following: NetBackup for SharePoint Server Administrator's Guide NetBackup for Exchange Server Administrator's Guide NetBackup for SQL Server Administrator's Guide
Actions > Edit	Edits the application host or component host of the currently selected mapping.
Actions > Delete	Deletes the mapping.

Encryption properties

To access these settings, in the web UI click **Hosts > Host properties**. Select the client. If necessary, click **Connect**, then click **Edit client**. Click **Encryption**.

The **Encryption** properties control encryption on the currently selected client.

More information is available in the [NetBackup Security and Encryption Guide](#).

The **Encryption permissions** property indicates the encryption setting on the selected NetBackup client as determined by the primary server.

Table 7-20 Encryption permissions selections

Property	Description
Not allowed	Specifies that the client does not permit encrypted backups. If the server requests an encrypted backup, the backup job ends due to error.
Allowed	Specifies that the client allows either encrypted or unencrypted backups. Allowed is the default setting for a client that has not been configured for encryption.
Required	Specifies that the client requires encrypted backups. If the server requests an unencrypted backup, the backup job ends due to error.

Choose the encryption properties.

Table 7-21 Encryption properties

Property	Description
Use standard encryption	Pertains to the 128-bit and the 256-bit options of NetBackup Encryption.
Client cipher	The following cipher types are available: AES-256-CFB and AES-128-CFB. AES-128-CFB is the default. Note: If you have 9.1 or earlier hosts in your environment, it is recommended that you select stronger client ciphers for the hosts, such as AES-256-CFB or AES-128-CFB. More information about the ciphers file is available in the NetBackup Security and Encryption Guide.

Additional encryption methods for Windows clients

In addition to NetBackup client and server data encryption, Microsoft Windows clients also have access to methods of encrypting the data on the original disk.

Each of the following methods has its own costs and benefits. NetBackup supports each method for protecting Microsoft Windows clients.

Encrypting File System

The Encrypting File System (EFS) on Microsoft Windows provides file system-level encryption. EFS is a form of encryption where individual files or directories are encrypted by the file system itself.

The technology enables files to be transparently encrypted to protect confidential data from attackers with physical access to the computer. Users can enable encryption on a per-file, per-directory, or per-drive basis. The Group Policy in a Windows domain environment can also mandate some EFS settings.

No NetBackup settings are involved in protecting these encrypted objects. Any object with an encrypted file system attribute is automatically backed up and restored in its encrypted state.

BitLocker Drive Encryption

BitLocker Drive Encryption is a full disk encryption feature included with Microsoft's Windows desktop and server versions.

Disk encryption is a technology which protects information by converting it into unreadable code that cannot be deciphered easily by unauthorized people. Disk encryption uses disk encryption software or hardware to encrypt every bit of data that goes on a disk or a disk volume.

As with EFS, no NetBackup settings are involved to use BitLocker for encryption. Unlike EFS, the encryption layer is invisible to NetBackup, with the data being automatically decrypted and encrypted by the operating system.

NetBackup does nothing to manage the encryption process and therefore backs up and restores the unencrypted data.

Note: If you recover a Windows computer that has BitLocker encryption enabled, you must re-enable BitLocker encryption following the restore.

Off-host backup is not supported with volumes that run Windows BitLocker Drive Encryption.

Enterprise Vault properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the Windows client. If necessary click **Connect**, then click **Edit client**. Click **Windows Client > Enterprise Vault**.

The **Enterprise Vault** properties apply to currently selected client .

To perform backups and restores, NetBackup must know the user name and password for the account that is used to log on to the Enterprise Vault Server and to interact with the Enterprise Vault SQL database. The user must set the logon account for every NetBackup client that runs backup and restore operations for Enterprise Vault components.

The **Enterprise Vault** host properties contains the following settings.

Table 7-22 Enterprise Vault properties

Property	Description
User name	Specify the user ID for the account that is used to log on to Enterprise Vault (DOMAIN\user name). Note: In 10.0 and later, credentials are stored in the Credential Management System (CMS).
Password	Specify the password for the account.
Consistency check before backup	Select what kind of consistency checks to perform on the SQL Server databases before NetBackup begins a backup operation.
Continue with backup if consistency check fails	Continues the backup job even if the consistency check fails. It may be preferable for the job to continue, even if the consistency check fails. For example, a backup of the database in its current state may be better than no backup at all. Or, it may be preferable for the backup of a large database to continue if it encounters only a small problem.

Enterprise Vault hosts properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Enterprise Vault hosts**.

The **Enterprise Vault hosts** properties apply to currently selected primary server.

Special configuration is required to allow NetBackup to restore SQL databases to the correct hosts in an Enterprise Vault farm. In the **Enterprise Vault hosts** primary server properties, specify a source and a destination host. By doing so, you specify a source host that can run restores on the destination host.

The **Enterprise Vault hosts** page contains the following properties.

Table 7-23 Enterprise Vault Hosts properties

Option	Description
Add	Adds the source and the destination hosts within the Enterprise Vault configuration. You must provide the name of the Source host and the name of the Destination host .
Actions > Edit	Changes the source host and the destination host.
Actions > Delete	Deletes the entry.

Exchange properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the Windows client. If necessary click **Connect**, then click **Edit client**. Click **Windows client > Exchange**.

The **Exchange** properties apply to the currently selected Windows client . For clustered or replicated environments, configure the same settings for all nodes. If you change the attributes for the virtual server name, only the DAG host server is updated.

For complete information on these options, see the [NetBackup for Exchange Server Administrator's Guide](#).

The **Exchange** host properties contain the following settings.

Table 7-24 Exchange properties

Property	Description
Backup option for log files during full backups	Note: This property only applies to MS-Exchange-Server backup policies. Choose which logs to include with snapshot backups: ■ Back up only uncommitted log files (not recommended for replication environments) ■ Back up all log files (including committed log files)
Exchange granular proxy host	Note: This property applies when you duplicate or browse a backup that uses Granular Recovery Technology (GRT). You can specify a different Windows system to act as a proxy for the source client when you duplicate or browse a backup (with <code>bplist</code>) that uses GRT. Use a proxy if you do not want to affect the source client or if it is not available.
Truncate Exchange log files after successful Instant Recovery backup	Note: This property only applies to MS-Exchange-Server backup policies. Enable this option to delete transaction logs after a successful Instant Recovery backup. By default, transaction logs are not deleted for a full Instant Recovery backup that is snapshot only.
Perform consistency check before backup with Microsoft Volume Shadow Copy Service (VSS)	Disable this option if you do not want to perform a consistency check during a DAG backup. If you select Continue with backup if consistency check fails, NetBackup continues to perform the backup even if the consistency check fails.

Table 7-24 Exchange properties (*continued*)

Property	Description
Exchange credentials	Note the following for this property: ■ This property applies to MS-Exchange-Server and VMware backup policies with Exchange recovery. ■ You must configure this property if you want to use GRT. Provide the credentials for the account for NetBackup Exchange operations. This account must have the necessary permissions to perform Exchange restores. The permissions that are required depend on the Exchange version that you have. The account also needs the right to "Replace a process level token."

About the Exchange credentials in the client host properties

The Exchange credentials in the client host properties indicate the account that has necessary permissions to perform Exchange restores. The permissions that are required depend on the Exchange version that you have.

Note the following:

- In NetBackup 10.0 and later, credentials are stored in the Credential Management System (CMS).
- To use GRT, configure the Exchange credentials on all granular clients. Alternatively, you can configure the Exchange credentials only on the granular clients that perform restores. In this case, for the entire domain add "Exchange Servers" to the "View-Only Organization Management" role group. Perform this configuration in the Exchange Administration Center (EAC) or in Active Directory. See the following Microsoft article for more information:
<http://technet.microsoft.com/en-us/library/jj657492>
- The account that you configured for the **Exchange credentials** must also have the right to "Replace a process level token."
- For database restores from VMware backups, the Exchange credentials that you provide must have permissions to restore VM files.
- If you want to restore from a VMware snapshot copy that was created with Replication Director, do the following:
 - Provide the Exchange credentials in the **Domain\user** and **Password** fields.
 - Configure the NetBackup Client Service with an account that has access to the CIFS shares that are created on the NetApp disk array.
- If you specify the minimal NetBackup account for the Exchange credentials in the client host properties, NetBackup can back up only active copies of the

Exchange databases. If you select **Passive copy only** in the **Exchange database backup source** field when you create a policy, any backups fail. The failure occurs because the Microsoft Active Directory Service Interface does not provide a list of database copies for a minimal account.

Exclude list properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the Windows client. If necessary click **Connect**, then click **Edit client**. Click **Windows client > Exclude list**.

Use the **Exclude list** host properties to create and to modify the exclude list for a Windows client . An exclude list names the files and directories to be excluded from backups.

If more than one exclude or include list exists for a client, NetBackup uses only the most specific one.

For example, assume that a client has the following exclude list:

- An exclude list for a policy and schedule.
- An exclude list for a policy.
- An exclude list for the entire client. This list does not specify a policy or schedule.

In this example, NetBackup uses the first exclude list (for policy and schedule) because it is the most specific.

Exclude and include lists do not determine if an entire drive is excluded when NetBackup determines if a backup job should start.

Normally, a problem does not occur. However, if a policy uses multistreaming and a drive or a mount point is excluded, that job reports an error status when it completes. To avoid the situation, do not use the policy or the policy and the schedule lists to exclude an entire volume.

The **Exclude list** host properties contain the following settings.

Table 7-25 Exclude list properties

Property	Description
Exclude list	Displays the excluded files and directories and the policies and schedules that they apply to.
Use case-sensitive exclude list	Indicates that the files and directories to exclude are case-sensitive.

Table 7-25 Exclude list properties (*continued*)

Property	Description
Exceptions to the exclude list	Displays any exceptions to the exclude list and the policies and schedules that they apply to. When the policies in this list run, the files and directories in the Exceptions to the exclude list are backed up. Adding an exception can be useful to exclude all files in a directory except one file. See “Add an exception to the exclude list” on page 134. For example, if the file list of items to back up contains <code>/foo</code>, and the exclude list contains <code>/foo/bar</code>, adding <code>/fum</code> to the exceptions list does not back up the <code>/fum</code> directory. However, adding <code>fum</code> to the exceptions list backs up any occurrences of <code>fum</code> (file or directory) that occur within <code>/foo/bar</code>.

Add an entry to an exclude list

Use the following procedure to add an entry to an exclude list for a policy or all policies. When the policies in the exclude list are run, the files and directories that are specified in the list are not backed up.

To add an entry to the exclude list

- 1 Open the NetBackup web UI.
- 2 On the left click **Hosts > Host properties**.
- 3 Select the client.
- 4 If necessary, click **Connect**. Then click **Edit client**.
- 5 Click **Windows clients > Exclude list**.
- 6 Under the Exclude list, click **Add**.
- 7 By default, the file, directory, or path are excluded from **All policies**. Or, type the name of the policy to exclude the items from a specific policy.
- 8 By default, the file, directory, or path are excluded from **All schedules**. Or, type the name of the schedule to exclude the items from a specific policy schedule.
- 9 Enter the file name, directory, or path that you want to exclude from the backups.
- 10 Click **Add**.

Add an exception to the exclude list

Use the following procedure to add an exception to the exclude list for a policy:

To add an exception to the exclude list

- 1** Open the NetBackup web UI.
- 2** On the left click **Hosts > Host properties**.
- 3** Select the client.
- 4** If necessary, click **Connect**. Then click **Edit client**.
- 5** Click **Windows clients > Exclude list**.
- 6** Expand **Exceptions to the exclude list**. Then click **Add**.
- 7** By default, the file, directory, or path is an exception for **All policies**. Or, type the name of the policy to add an exception for a specific policy.
- 8** By default, the file, directory, or path for **All schedules**. Or, type the name of the schedule to add an exception for a specific policy schedule.
- 9** Enter the file name, directory, or path that you want to exclude from the backups.
- 10** Click **Add**.

Syntax rules for exclude lists

It is recommended that you always specify automounted directories and CD-ROM file systems in the exclude list. Otherwise, if the directories are not mounted at the time of a backup, NetBackup must wait for a timeout.

The following syntax rules apply to exclude lists:

- Only one pattern per line is allowed.
- NetBackup recognizes standard wildcard use.
See [“Wildcard use in NetBackup”](#) on page 615.
See [“NetBackup naming conventions”](#) on page 614.
- If all files are excluded in the backup selections list, NetBackup backs up only what is specified by full path names in the include list. Files can be excluded by using / or * or by using both symbols together (/*).
- Spaces are considered legal characters. Do not include extra spaces unless they are part of the file name.
For example, if you want to exclude a file named
`C:\testfile` (with no extra space character at the end)
and your exclude list entry is
`C:\testfile` (with an extra space character at the end)
NetBackup cannot find the file until you delete the extra space from the end of the file name.

- End a file path with `\` to exclude only directories with that path name (for example, `C:\users\test\`). If the pattern does not end in `\` (for example, `C:\users\test`), NetBackup excludes both files and directories with that path name.
- To exclude all files with a given name, regardless of their directory path, enter the name. For example:

`test`

rather than

`C:\test`

This example is equivalent to prefixing the file pattern with

`\`

`\*\`

`\*\*\`

`\*\*\*\`

and so on.

The following syntax rules apply only to UNIX clients:

- Do not use patterns with links in the names. For example, assume `/home` is a link to `/usr/home` and `/home/doc` is in the exclude list. The file is still backed up in this case because the actual directory path, `/usr/home/doc`, does not match the exclude list entry, `/home/doc`.
- Blank lines or lines which begin with a pound sign (`#`) are ignored.

Example of a Windows client exclude list

Assume that an exclude list in the **Exclude list** host properties contains the following entries:

`C:\users\doe\john`

`C:\users\doe\abc\`

`C:\users\*\test`

`C:\*\temp`

`core`

Given the exclude list example, the following files, and directories are excluded from automatic backups:

- The file or directory named `C:\users\doe\john`.
- The directory `C:\users\doe\abc\` (because the exclude entry ends with `\`).
- All files or directories named `test` that are two levels beneath `users` on drive C.

- All files or directories named `temp` that are two levels beneath the root directory on drive C.
- All files or directories named `core` at any level and on any drive.

Example of a UNIX exclude list

In this example of a UNIX exclude list, the list contains the following entries:

```
# this is a comment line
/home/doe/john
/home/doe/abc/
/home/*/test
/*temp
core
```

Given the exclude list example, the following files and directories are excluded from automatic backups:

- The file or directory named `/home/doe/john`.
- The directory `/home/doe/abc` (because the exclude entry ends with `/`).
- All files or directories named `test` that are two levels beneath `home`.
- All files or directories named `temp` that are two levels beneath the root directory.
- All files or directories named `core` at any level.

About creating an include list on a UNIX client

To add a file that is eliminated with the exclude list, create a `/usr/opensv/netbackup/include_list` file. The same syntax rules apply as for the exclude list.

Note: Exclude and include lists do not apply to user backups and archives.

To illustrate the use of an include list, we use the example from the previous discussion. The exclude list in that example causes NetBackup to omit all files or directories named `test` from all directories beneath `/home/*/test`.

In this case, add a file named `/home/jdoe/test` back into the backup by creating an `include_list` file on the client. Add the following to the `include_list` file:

```
# this is a comment line
/home/jdoe/test
```

To create an include list for a specific policy or policy and schedule combination, use a `.policyname` or `.policyname.schedulename` suffix. The following are two examples of include list names for a policy that is named `wkstations` that contains a schedule that is named `fulls`.

```
/usr/opensv/netbackup/include_list.workstations  
/usr/opensv/netbackup/include_list.workstations.fulls
```

The first file affects all scheduled backups in the policy that is named `wkstations`. The second file affects backups only when the schedule is named `fulls`.

For a given backup, NetBackup uses only one include list: the list with the most specific name. Given the following two files:

```
include_list.workstations  
include_list.workstations.fulls
```

NetBackup uses only `include_list.workstations.fulls` as the include list.

Traversing excluded directories

An exclude list can indicate a directory for exclusion, while the client uses an include list to override the exclude list. NetBackup traverses the excluded directories if necessary, to satisfy the client's include list.

Assume the following settings for a Windows client:

- The backup policy backup selection list indicates `ALL_LOCAL_DRIVES`. When a scheduled backup runs, the entire client is backed up.
The entire client is also backed up if the backup selection list consists of only:
/
 - The exclude list on the client consists of only: *An exclude list of * indicates that all files are excluded from the backup.
- However, since the include list on the Windows client includes the following file:
`C:\WINNT`, the excluded directories are traversed to back up `C:\WINNT`.
If the include list did not contain any entry, no directories are traversed.

In another example, assume the following settings for a UNIX client:

- The backup selection list for the client consists of the following: /
- The exclude list for the UNIX client consists of the following: /
- The include list of the UNIX client consists of the following directories:
/data1
/data2
/data3

Because the include list specifies full paths and the exclude list excludes everything, NetBackup replaces the backup selection list with the client's include list.

Fibre transport properties

NetBackup Fibre Transport properties control how your Fibre Transport media servers and SAN clients use the Fibre Transport service for backups and restores. The **Fibre transport** properties apply to the host type that you select, as follows:

Table 7-26 Host types for Fibre transport properties

Host type	Description
Primary server	Global Fibre transport properties that apply to all SAN clients.
Media server	The Fibre transport Maximum concurrent FT connections property applies to the FT media server that you select.
Client	The Fibre transport properties apply to the SAN client that you select. The default values for clients are the global property settings of the primary server. Client properties override the global Fibre transport properties.

The **Fibre transport** properties contain the following settings. All properties are not available for all hosts. In this table, FT device is an HBA port on a Fibre Transport media server. The port carries the backup and restore traffic. A media server may have more than one FT device.

Table 7-27 Fibre transport properties

Property	Description
Maximum concurrent FT connections	This property appears only when you select an FT media server . This property specifies the number of FT connections to allow to the selected media server or media servers. A connection is equivalent to a job. If no value is set, NetBackup uses the following defaults: ■ For NetBackup Appliance model 5330 and later: 32 ■ For NetBackup Appliance model 5230 and later: 32 ■ For NetBackup Fibre Transport media servers: 8 times the number of fast HBA ports on the media server plus 4 times the number of slow HBA ports. A fast port is 8 GB or faster, and a slow port is less than 8 GB. You can enter up to the following maximum connections for the media server or servers to use: ■ On a Linux FT media server host: 40. It is recommended that you use 32 or fewer connections concurrently on Linux. On Linux hosts, you can increase that maximum by setting a NetBackup touch file, <code>NUMBER_DATA_BUFFERS_FT</code>. See “About Linux concurrent FT connections” on page 141. ■ For NetBackup Appliance model 5330 and later: 40. ■ For NetBackup Appliance model 5230 and later: 40. ■ On a Solaris FT media server host: 64. NetBackup supports 644 buffers per media server for Fibre Transport. To determine the number of buffers that each connection uses, divide 644 by the value you enter. More buffers per connection equal better performance for each connection.
Use defaults from the primary server configuration	This property appears only when you select a client . This property specifies that the client follow the properties as they are configured on the primary server.
Preferred	The Preferred property specifies to use an FT device if one is available within the configured wait period in minutes. If an FT device is not available after the wait period elapses, NetBackup uses a LAN connection for the operation. If you select this option, also specify the wait period for backups and for restores. For the global property that is specified on the primary server, the default is Preferred.

Table 7-27 Fibre transport properties (*continued*)

Property	Description
Always	The Always property specifies that NetBackup should always use an FT device for backups and restores of SAN clients. NetBackup waits until an FT device is available before it begins the operation. However, an FT device must be online and up. If not, NetBackup uses the LAN. An FT device may be unavailable because none are active, none have been configured, or the SAN Client license expired.
Fail	The Fail property specifies that NetBackup should fail the job if an FT device is not online and up. If the FT devices are online but busy, NetBackup waits until a device is available and assigns the next job to the device. An FT device may be unavailable because none are active, none have been configured, or the SAN Client license expired.
Never	The Never property specifies that NetBackup should never use an FT pipe for backups and restores of SAN clients. NetBackup uses a LAN connection for the backups and restores. If you specify Never for the primary server, Fibre Transport is disabled in the NetBackup environment. If you select Never, you can configure FT usage on a per-client basis. If you specify Never for a media server, Fibre Transport is disabled for the media server. If you specify Never for a SAN client, Fibre Transport is disabled for the client.

NetBackup provides one finer level of granularity for Fibre transport. SAN client usage preferences override the FT properties that you configure through **Host properties**.

About Linux concurrent FT connections

NetBackup uses the **Maximum concurrent FT connections** setting in the **Fibre transport** host property to configure the number of concurrent connections to a Fibre transport media server, up to the total that is allowed per host.

See “[Fibre transport properties](#)” on page 139.

If the total number of concurrent connections on Linux is too low for your purposes, you can increase the total number of concurrent connections. The consequence is that each client backup or restore job uses fewer buffers, which means that each job is slower because of fewer buffers. To increase the number of concurrent connections, reduce the number of buffers per connection. To do so, create the following file and include one of the supported values from [Table 7-28](#) in the file:

/usr/opensv/netbackup/db/config/NUMBER_DATA_BUFFERS_FT

Table 7-28 shows the values that NetBackup supports for the NUMBER_DATA_BUFFERS_FT file. NetBackup supports 644 buffers per media server for Fibre transport.

Table 7-28 Supported values for buffers per FT connection

NUMBER_DATA_BUFFERS_FT	Total concurrent connections: NetBackup 5230 and 5330 and later appliances	Total concurrent connections: Linux FT media server
16	40	40
12	53	53
10	64	64

If you want, you then can limit the number of connections for a media server with the **Maximum concurrent FT connections** setting in the **Fibre transport** host properties.

Firewall properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server or media server. If necessary click **Connect**, then click **Edit primary server** or **Edit media server**. Click **Firewall**.

The **Firewall** properties determine how the selected primary servers and media servers connect to the legacy services that run on that NetBackup host.

Servers are added to the **Hosts** list of the **Firewall** properties. To configure port usage for clients, see the **Client attributes** properties.

See “[Client attributes properties](#)” on page 104.

The **Firewall** host properties contain the following settings.

Table 7-29 Firewall properties

Property	Description
Default connect options	By default, the Default connect options include firewall-friendly connect options including the fewest possible ports to open. The default options can be set differently for an individual server or client with the settings in Attributes for selected hosts. To change the default connect options for the selected server or client, click Edit. These properties correspond to the <code>DEFAULT_CONNECT_OPTIONS</code> configuration option.
Hosts	You can configure different default connect options for the hosts that are displayed in this list. - Click Add to add a host to the Hosts list. You must add a host name to the list before you can configure different settings for that host. Servers do not automatically appear on the list. - To configure different settings for a host, select the host name in the Hosts list. Then select the connect options in the Attributes for selected hosts section. - To remove the host from the list, locate a host name in the list. Then click Delete.
Attributes for selected hosts	This section displays the connect options for the selected server. To change the connection options for a server, first select the host name in the Hosts list. These properties correspond to the <code>CONNECT_OPTIONS</code> configuration option.
BPCD connect back	This property specifies how daemons are to connect back to the NetBackup Client daemon (<code>BPCD</code>) as follows: Use default connect options (An option for individual hosts) Use the methods that are specified under Default connect options. Random port NetBackup randomly chooses a free port in the allowed range to perform the traditional connect-back method. VNETD port This method requires no connect-back. The Cohesity Network Daemon (<code>vnetd</code>) was designed to enhance firewall efficiency with NetBackup during server-to-server and server-to-client communications. The server initiates all <code>bpcd</code> socket connections. Consider the example in which <code>bpbrm</code> on a media server initially connects with <code>bpcd</code> on a client. The situation does not pose a firewall problem because <code>bpbrm</code> uses the well-known PBX or <code>vnetd</code> port.

Table 7-29 Firewall properties (*continued*)

Property	Description
Ports	Select whether a reserved or non-reserved port number should be used to connect to the host name: ■ Use default connect options (An option for individual hosts) Use the methods that are specified under Default attributes. ■ Reserved ports Connect to the host name by a reserved port number. ■ Non-reserved ports Connect to the host name by a non-reserved port number. To configure port usage for clients, see the Client attributes properties.

General server properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server or media server. If necessary click **Connect**, then click **Edit primary server** or **Edit media server**. Click **General server**.

The **General server** properties apply to the selected primary server or media server.

The **General server** page contains the following properties.

Table 7-30 General server properties

Property	Description
Delay on multiplexed restores	This property specifies how long the server waits for additional restore requests of multiplexed images on the same tape. All of the restore requests that are received within the delay period are included in the same restore operation (one pass of the tape). The default is a delay of 30 seconds.
Check the capacity of disk storage units every	This property applies to the disk storage units of 6.0 media servers only. Subsequent releases use internal methods to monitor disk space more frequently.

Table 7-30 General server properties (*continued*)

Property	Description
Must use local drive	This property appears for primary servers only, but applies to all media servers as well. This property does not apply to NDMP drives. If a client is also a media server or a primary server and Must use local drive is selected, a local drive is used to back up the client. If all drives are down, another can be used. This property increases performance because backups are done locally rather than sent across the network. For example, in a SAN environment a storage unit can be created for each SAN media server. Then, the media server clients may be mixed with other clients in a policy that uses ANY AVAILABLE storage unit. When a backup starts for a client that is a SAN media server, the backups go to the SAN connected drives on that server.
Use direct access recovery for NDMP restores	By default, NetBackup for NDMP is configured to use Direct Access Recovery (DAR) during NDMP restores. DAR can reduce the time it takes to restore files by allowing the NDMP host to position the tape to the exact location of the requested file(s). Only the data that is needed for those files is read. Clear this check box to disable DAR on all NDMP restores. Without DAR, NetBackup reads the entire backup image, even if only a single restore file is needed.
Enable message-level cataloging when duplicating Exchange images that use Granular Recovery Technology	This option performs message-level cataloging when you duplicate Exchange backup images that use Granular Recovery Technology (GRT) from disk to tape. To perform duplication more quickly, you can disable this option. However, then users are not able to browse for individual items on the image that was duplicated to tape. See the NetBackup for Exchange Administrator's Guide.

Table 7-30 General server properties (*continued*)

Property	Description
Media host override list	Specific servers can be specified in this list as servers to perform restores, regardless of where the files were backed up. (Both servers must be in the same primary and media server cluster.) For example, if files were backed up on media server A, a restore request can be forced to use media server B. The following items describe situations in which the capability to specify servers is useful: ■ Two (or more) servers share a robot and each have connected drives. A restore is requested while one of the servers is either temporarily unavailable or is busy doing backups. ■ A media server was removed from the NetBackup configuration, and is no longer available. To add a host to the Media host override list, click Add. To change an entry in the list, select a host name, then click Actions > Edit. Configure the following options: ■ Original backup server Enter the name of the server where the data was backed up originally. ■ Restore server Enter the name of the server that is to process future restore requests.

Forcing restores to use a specific server

Use the following procedure to force restores to use a specific server.

To force restores to use a specific server

- 1 If necessary, physically move the media to the host to answer the restore requests, then update the NetBackup database to reflect the move.
- 2 Modify the NetBackup configuration on the primary server.
 - Open the NetBackup web UI and sign into the primary server.
 - On the left, click **Host > Host properties**.
 - Select the primary server.
 - If necessary, click **Connect**. Then click **Edit primary server**.
 - Click **General server**.

- Add the original backup media server and the restore server to the **Media host override** list.
- 3 Stop and restart the NetBackup Request Daemon (`bprd`) on the primary server.
- This process applies to all storage units on the original backup server. Restores for any storage unit on the **Original backup server** go to the server that is listed as the **Restore server**.
- To revert to the original configuration for future restores, delete the line from the **Media host override** list.

Global attributes properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Global attributes**.

The **Global attributes** properties apply to currently selected primary servers. These properties affect all operations for all policies and clients. The default values are adequate for most installations.

The **Global attributes** page contains the following properties.

Table 7-31 Global attributes properties

Property	Description
Job retry delay	This property specifies how often NetBackup retries a job. The default is 10 minutes. The maximum is 60 minutes; the minimum is 1 minute.

Table 7-31 Global attributes properties (*continued*)

Property	Description
Maximum jobs per second	This property specifies the throttle on the maximum number of backup jobs that are allowed to go from the Queued to Active state per second. By default, the value of this property is 0, which means no throttling occurs. After the maximum number of jobs is reached in one second, subsequent jobs will remain in the Queued state. In the next second, jobs are released in a first-in-first-out order from the Queued state until the maximum jobs value is reached again or until all throttled jobs or new jobs have been made active. This property can be used to smooth out the resource utilization curve. It is particularly useful when backup windows open and a large number of jobs are scheduled to start within a short time period. This value supersedes the <code>DBM_NEW_IMAGE_DELAY</code> configuration value found here: https://www.veritas.com/support/en_US/article.100047119 If <code>DBM_NEW_IMAGE_DELAY</code> is configured and the maximum jobs per second throttle is the default value, the <code>DBM_NEW_IMAGE_DELAY</code> will be converted to an equivalent jobs-per-second value. It will not modify the configuration. For example, if <code>DBM_NEW_IMAGE_DELAY</code> was set to 333ms, the NetBackup Job Manager will use a maximum jobs per second throttle of 3. If the user then configured the maximum jobs throttle to 2 per second, the configured <code>DBM_NEW_IMAGE_DELAY</code> would be ignored. Note: This throttle only affects the number of backup jobs that the NetBackup Job Manager allows to start in one second. It does not affect other job types like restores, archives, duplications, or replications. It does not affect the maximum number of concurrent jobs.
Maximum jobs per client	This property specifies the maximum number of backup and archive jobs that NetBackup clients can perform concurrently. The default is one job. NetBackup can process concurrent backup jobs from different policies on the same client only in the following situations: ■ More than one storage unit available ■ One of the available storage units can perform more than one backup at a time. See “About constraints on the number of concurrent jobs” on page 150.
Policy update interval	This property specifies how long NetBackup waits to process a policy after a policy is changed. The interval allows the NetBackup administrator time to make multiple changes to the policy. The default is 10 minutes. The maximum is 1440 minutes; the minimum is 1 minute.
Compress catalog interval	This property specifies how long NetBackup waits after a backup before it compresses the image catalog file.

Table 7-31 Global attributes properties (*continued*)

Property	Description
Schedule backup attempts	NetBackup considers the failure history of a policy to determine whether or not to run a scheduled backup job. The Schedule backup attempts property sets the timeframe for NetBackup to examine. This property determines the following characteristics for each policy: - How many preceding hours NetBackup examines to determine whether to allow another backup attempt (retry). By default, NetBackup examines the past 12 hours. - How many times a backup can be retried within that timeframe. By default, NetBackup allows two attempts. Attempts include the scheduled backups that start automatically or the scheduled backups that are user-initiated. Consider the following example scenario using the default setting 2 tries every 12 hours: - Policy_A runs at 6:00 P.M.; Schedule_1 fails. - Policy_A is user-initiated at 8:00 P.M.; Schedule_2 fails. - At 11:00 P.M., NetBackup looks at the previous 12 hours. NetBackup sees one attempt at 6:00 P.M. and one attempt at 8:00 P.M. The Schedule backup attempts setting of two has been met so NetBackup does not try again. - At 6:30 A.M. the next morning, NetBackup looks at the previous 12 hours. NetBackup sees only one attempt at 8:00 P.M. The Schedule backup attempts setting of two has not been met so NetBackup tries again. If a schedule window is not open at this time, NetBackup waits until a window is open. Note: This attribute does not apply to user backups and archives.
Maximum vault jobs	This property specifies the maximum number of vault jobs that are allowed to be active on the primary server. The greater the maximum number of vault jobs, the more system resources are used. If the active vault jobs limit is reached, subsequent vault jobs are queued and their status is shown as Queued in the Activity monitor. If a duplication job or eject job waits, its status is shown as Active in the Activity monitor. See “Job monitoring” on page 49.
Administrator email address property	This property specifies the addresses where NetBackup sends notifications of scheduled backups or administrator-directed manual backups. To send the information to more than one administrator, separate multiple email addresses by using a comma, as follows: <pre>useraccount1@company.com,useraccount2@company.com</pre> More information is available on the configuration requirements for email notifications. See “Send notifications to the backup administrator about failed backups” on page 72.

About constraints on the number of concurrent jobs

Specify any number of concurrent jobs within the following constraints.

Table 7-32 Constraints on concurrent jobs

Constraint	Description
Number of storage devices	NetBackup can perform concurrent backups to separate storage units or to drives within a storage unit. For example, a single Media Manager storage unit supports as many concurrent backups as it has drives. A disk storage unit is a directory on disk, so the maximum number of jobs depends on system capabilities.
Server and client speed	Too many concurrent backups on an individual client interfere with the performance of the client. The best setting depends on the hardware, operating system, and applications that are running. The Maximum jobs per client property applies to all clients in all policies. To accommodate weaker clients (ones that can handle only a small number of jobs concurrently), consider using one of the following approaches: ■ Set the Maximum data streams property for those weaker clients appropriately. (Open the host properties for the primary server. Then click Client attributes > General tab.) See “General tab of the Client attributes properties” on page 106. ■ Use the Limit jobs per policy policy setting in a client-specific policy. (A client-specific policy is one in which all clients share this characteristic).
Network loading	The available bandwidth of the network affects how many backups can occur concurrently. The load might be too much for a single Ethernet. For loading problems, consider backups over multiple networks or compression. A special case exists to back up a client that is also a server. Network loading is not a factor because the network is not used. Client and server loading, however, is still a factor.

Note: Catalog backups can run concurrently with other backups. To do so, set the **Maximum jobs per client** value to greater than two for the primary server. The higher setting ensures that the catalog backup can proceed while the regular backup activity occurs.

Setting up mailx email client

NetBackup supports setting up email notifications by using mailx client.

To set up a mailx email client

- 1 Navigate to the /etc/mail.rc location.
- 2 Edit the file to add the SMTP server settings.

For example, set

```
smtp=<Your_SMTP_Server_Hostname>:<SMTP_SERVER_PORT>
```

Logging properties

To access the Logging properties, in the web UI select **Hosts > Host properties**. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Logging**.

The logging settings determine the behavior for NetBackup logging on the primary server, media server, and the clients:

- Overall logging level or global logging level for all NetBackup processes.
- Overrides for the specific processes that use legacy logging.
- Logging levels for the services that use unified logging.
- Logging for critical processes.
- On clients, the logging level for database applications.
- Log retention settings for NetBackup and for NetBackup Vault (if it is installed).

All NetBackup processes use either unified logging or legacy logging. You can set a global or a unique logging level for certain processes and services. Retention levels limit the size of the log files or (for the primary server) the number of days the logs are kept. If you use NetBackup Vault, you can select separate logging retention settings for that option.

For complete details on logging, see the [NetBackup Logging Reference Guide](#).

Table 7-33 Logging properties

Property	Description
Global logging level	This setting establishes a global logging level for all processes that are set to Same as global. The Global logging level affects the legacy and unified logging level of all NetBackup processes on the server or client. This setting does not affect the following logging processes: ■ PBX logging See the NetBackup Troubleshooting Guide for more information on how to access the PBX logs. ■ Media and device management logging (vmd, ltid, avrd, robotic daemons, media manager commands)
Process-specific overrides	These settings let you override the logging level for the specific processes that use legacy logging.
Debug logging levels for NetBackup services	These settings let you manage the logging level for the specific services that use unified logging.
Logging for critical processes	The option lets you enable logging for the critical processes: ■ Primary server processes: bprd and bpbm. ■ Media server processes: bpbm, bptm, and bpdm. ■ Client process: bpfis Note the following: ■ If you enable Logging for critical processes, also enable the option Maximum log size. If you disable this option it may adversely affect NetBackup operations. ■ This option sets the log retention to the default log size. ■ Clicking Restore to defaults does not modify the Logging for critical processes or the Maximum log size options. ■ To disable the logging for critical processes, modify the logging levels for those processes.
Retention period	Specifies the length of time NetBackup keeps information from the error catalog, job catalog, and debug logs. Note that NetBackup derives its reports from the error catalog. The logs can consume a large amount of disk space, so do not keep the logs any longer than necessary. The default is 28 days. Note: This setting is not applicable for Cloud Scale.

Table 7-33 Logging properties (*continued*)

Property	Description
Maximum log size	Specifies the size of the NetBackup logs that you want to retain. When the NetBackup log size grows to this value, the older logs are deleted. ■ For primary and media servers, the recommended value is 25 GB or greater. ■ For clients, the recommended value is 5 GB or greater. Note: This setting is not applicable for Cloud Scale.
Vault logs retention period	If NetBackup Vault is installed, select the number of days to keep the Vault session directories, or select Forever .

Logging levels

You can choose to apply the same logging level for all NetBackup processes. Or, you can select logging levels for specific processes or services.

Table 7-34 Logging level descriptions

Logging level	Description
Same as global	The process uses the same logging level as the Global logging level .
No logging	No log is created for the process.
Minimum logging (default)	A small amount of information is logged for the process. Use this setting unless advised otherwise by Cohesity Technical Support. Other settings can cause the logs to accumulate large amounts of information.
Levels 1 through 4	Progressively more information is logged at each level for the process.
5 (Maximum)	The maximum amount of information is logged for the process.

Global logging level

This setting controls the logging level for all processes and for those processes that are set to **Same as global**. You can control the logging level for some NetBackup processes individually.

See [the section called “Overrides for legacy logging levels”](#) on page 154.

See [the section called “Unified logging levels for the primary server”](#) on page 154.

Overrides for legacy logging levels

These logging levels apply to legacy processes logging. The logging levels that are displayed depend on the type of host (primary, media, or client).

Table 7-35 Logging level overrides for legacy processes

Service	Description	Primary server	Media server	Client
BPBRM logging level	The NetBackup backup and restore manager.	X	X	
BPDM logging level	The NetBackup disk manager.	X	X	
BPTM logging level	The NetBackup tape manager.	X	X	
BPJOB logging level	The NetBackup Jobs Database Management daemon. This setting is only available for the primary server.	X		
BPDBM logging level	The NetBackup database manager.	X		
BPRD logging level	The NetBackup Request Daemon.	X		
Database logging level	The logging level for database agent logs. For details on which logs to create and refer to, see the guide for the specific agent.			X

Unified logging levels for the primary server

These logging levels apply to NetBackup services logging and are only available for the primary server.

Table 7-36 Logging levels for NetBackup services

Service	Description
Policy execution manager	The Policy execution manager (NBP) creates policy and client tasks and determines when jobs are due to run. If a policy is modified or if an image expires, NBP is notified and the appropriate policy and client tasks are updated.
Job manager	The Job Manager (NBJM) accepts the jobs that the Policy Execution Manager submits and acquires the necessary resources.
Resource broker	The Resource Broker (NBRB) makes the allocations for storage units, tape drives, client reservations.

Logging values in the registry, bp.conf file, and unified logging

You can also set logging values in the Windows registry, the bp.conf file, or in unified logging.

Table 7-37 Logging levels and their values

Logging level	Legacy logging - Windows registry	Legacy logging - bp.conf	Unified logging
Minimum logging	Hexadecimal value of 0xffffffff.	VERBOSE = 0 (global) processname_VERBOSE = 0 If the global VERBOSE value is set to a value other than 0, an individual process can be decreased by using the value -1. For example, processname_VERBOSE = -1.	1
No logging	Hexadecimal value of 0xfffffffffe.	VERBOSE=-2 (global) processname_VERBOSE = -2	0

Lotus Notes properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the client and click **Edit client**. Then click **Windows clients > Lotus Notes** or **UNIX client > Lotus Notes**.

The **Lotus Notes** properties apply to the currently selected client that runs NetBackup for Domino.

For more information, see the [NetBackup for HCL Domino Administrator's Guide](#).

For UNIX servers: If you have multiple installations of Domino server, the values in the client properties only apply to one installation. For other installations, specify the installation path and location of the `notes.ini` file with the `LOTUS_INSTALL_PATH` and `NOTES_INI_PATH` directives in the backup policy.

Table 7-38 Lotus Notes client host properties

Client host properties	Description
Maximum number of logs to restore	The maximum number of logs that can be prefetched in a single restore job during recovery. Specify a value greater than 1. A value less than or equal to 1, does not gather transaction logs during recovery. One transaction log extent per job is restored to the Domino server's log directory.
Transaction log cache path	A path where NetBackup can temporarily store the prefetched transaction logs during recovery. If you do not specify a path, during recovery NetBackup restores the logs to the Domino server's transaction log directory. Note the following: ■ If the specified path does not exist then it is created during restore. ■ The user must have write permission for the folder. ■ Transaction logs are restored to the original location, the Domino transaction log directory, if a path is not specified. ■ If the value of Maximum number of logs to restore is less than or equal to 1 then this path is ignored. The logs are not prefetched; one transaction log per job is restored to the Domino Server's log directory. ■ If there is not sufficient space to restore the specified number of logs, NetBackup tries to restore only the number of logs that can be accommodated.
INI path	The <code>notes.ini</code> file that is associated with the Domino partitioned servers used to back up and restore the Notes database. This setting does not apply to non-partitioned servers. ■ On Windows: If the <code>notes.ini</code> file is not located in the default directory, indicate its location. ■ On UNIX: If the <code>notes.ini</code> is not located in the directory that is specified in the Path, indicate its location here. Include the directory and the <code>notes.ini</code> file name.
Path	The path where the Notes program files reside on the client. NetBackup must know where these files are to perform backup and restore operations. ■ On Windows: The path for program directory (where <code>nserver.exe</code> resides). ■ On UNIX: A path that includes the Domino data directory, the Notes program directory, and the Notes resource directory.

Media properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server. If necessary, click **Connect**. Then click **Edit primary server** or **Edit media server**. Click **Media**.

The **Media** host properties contain the following settings.

Table 7-39 Media properties

Property	Description
Allow media overwrite property	This property overrides the NetBackup overwrite protection for specific media types. Normally, NetBackup does not overwrite certain media types. To disable overwrite protection, place a check in the check box of one or more of the listed media formats. For example, place a check in the CPIO check box to permit NetBackup to overwrite the cpio format. By default, NetBackup does not overwrite any of the formats on removable media, and logs an error if an overwrite attempt occurs. This format recognition requires that the first variable length block on a media be less than or equal to 32 kilobytes. The following media formats on removable media can be selected to be overwritten: ■ When ANSI is enabled, ANSI labeled media can be overwritten. ■ When TAR is enabled, TAR media can be overwritten. ■ When DBR is enabled, DBR media can be overwritten. (The DBR backup format is no longer used.) ■ Remote Storage MTF1 media format. When MTF1 is enabled, Remote Storage MTF1 media format can be overwritten. ■ When CPIO is enabled, CPIO media can be overwritten. ■ When AOS/VS is enabled, AOS/VS media can be overwritten. (Data General AOS/VS backup format.) ■ When MTF is enabled, MTF media can be overwritten. With only MTF checked, all other MTF formats can be overwritten. (The exception is Backup Exec MTF (BE-MTF1) and Remote Storage MTF (RS-MTF1) media formats, which are not overwritten. ■ When BE-MTF1 is enabled, Backup Exec MTF media can be overwritten. See “Results when media overwrites are not permitted” on page 160.

Table 7-39 Media properties (*continued*)

Property	Description
Enable SCSI reserve	This property allows exclusive access protection for tape drives. With access protection, other host bus adaptors cannot issue commands to control the drives during the reservation. SCSI reservations provide protection for NetBackup Shared Storage Option environments or any other multiple-initiator environment in which drives are shared. The protection setting configures access protection for all tape drives from the media server on which the option is configured. You can override the media server setting for any drive path from that media server. See “Recommended use for Enable SCSI reserve property” on page 161. The following are the protection options: ■ The SCSI persistent reserve option provides SCSI persistent reserve protection for SCSI devices. The devices must conform to the SCSI Primary Commands - 3 (SPC-3) standard. ■ The SPC-2 SCSI reserve option (default) provides SPC-2 SCSI reserve protection for SCSI devices. The devices must conform to the reserve and release management method in the SCSI Primary Commands - 2 standard. ■ To operate NetBackup without tape drive access protection, clear the Enable SCSI reserve property. If unchecked, other HBAs can send the commands that may cause a loss of data to tape drives. Note: Ensure that all of your hardware processes SCSI persistent reserve commands correctly. All of your hardware includes Fibre Channel bridges. If the hardware does not process SCSI persistent reserve commands correctly and NetBackup is configured to use SCSI persistent reserve, no protection may exist.
Allow multiple retentions per media	This property lets NetBackup mix retention levels on tape volumes. It applies to media in both robotic drives and nonrobotic drives. The default is that the check box is clear and each volume can contain backups of only a single retention level.
Allow backups to span tape media	This property, when checked, lets backups span to multiple tape media. This property lets NetBackup select another volume to begin the next fragment. The resulting backup has data fragments on more than one volume. The default is that Allow backups to span tape media is checked and backups are allowed to span media. If the end of media is encountered and this property is not selected, the media is set to FULL and the operation terminates abnormally. This action applies to both robotic drives and nonrobotic drives.

Table 7-39 Media properties (*continued*)

Property	Description
Allow backups to span disk volumes	This property lets backups span disk volumes when one disk volume becomes full. The default is that this property is enabled. The Allow backups to span disk volumes property does not apply to AdvancedDisk or OpenStorage storage units. Backups span disk volumes within disk pools automatically. The following destinations support disk spanning: ■ A BasicDisk storage unit spanning to a BasicDisk storage unit. The units must be within a storage unit group. ■ An OpenStorage or AdvancedDisk volume spanning to another volume in the disk pool. For disk spanning to occur, the following conditions must be met: ■ The storage units must share the same media server. ■ The multiplexing level on spanning storage units should be the same. If there are any differences, the level on the target unit can be higher. ■ A disk staging storage unit cannot span to another storage unit. Also, a disk staging storage unit is not eligible as a target for disk spanning. ■ Disk spanning is not supported on NFS.
Enable standalone drive extension	This property lets NetBackup use whatever labeled or unlabeled media is found in a nonrobotic drive. The default is that the Enable standalone drive extension property is enabled.
Enable job logging	This property allows the logging of the job information. This logging is the same information that the NetBackup Activity monitor uses. The default is that job logging occurs.
Enable unrestricted media sharing for all media servers	This property controls media sharing, as follows: ■ Enable this property to allow all NetBackup media servers and NDMP hosts in the NetBackup environment to share media for writing. Do not configure server groups for media sharing. ■ Clear this property to restrict media sharing to specific server groups. Then configure media server groups and backup policies to use media sharing. ■ Clear this property to disable media sharing. Do not configure media server groups. The default is that media sharing is disabled. (The property is cleared and no server groups are configured.) See “About NetBackup server groups” on page 263.

Table 7-39 Media properties (*continued*)

Property	Description
Media ID prefix (non-robotic)	This property specifies the media ID prefix to use in media IDs when the unlabeled media is in nonrobotic drives. The prefix must be one to three alpha-numeric characters. NetBackup appends numeric characters. By default, NetBackup uses A and assigns media IDs such as A00000, A00001, and so on. For example, if FEB is specified, NetBackup appends the remaining numeric characters. The assigned media IDs become FEB000, FEB001, and so on.
Media unmount delay	To specify a Media unmount delay property indicates that the unloading of media is delayed after the requested operation is complete. Media unmount delay applies only to user operations, to include backups and restores of database agent clients, such as those running NetBackup for Oracle. The delay reduces unnecessary media unmounts and the positioning of media in cases where the media is requested again a short time later. The delay can range from 0 seconds to 1800 seconds. The default is 180 seconds. If you specify 0, the media unmount occurs immediately upon completion of the requested operation. Values greater than 1800 are set to 1800.
Media request delay (non-robotic)	This property specifies how long NetBackup waits for media in nonrobotic drives. During the delay period, NetBackup checks every 60 seconds to see if the drive is ready. If the drive is ready, NetBackup uses it. Otherwise, NetBackup waits another 60 seconds and checks again. If the total delay is not a multiple of 60, the last wait is the remainder. If the delay is less than 60 seconds, NetBackup checks after the end of the delay. For example, set the delay to 150 seconds. NetBackup waits 60 seconds, checks for ready, waits 60 seconds, checks for ready, waits 30 seconds, and checks for ready the last time. If the delay was 50 seconds (a short delay is not recommended), NetBackup checks after 50 seconds.

Results when media overwrites are not permitted

If media contains one of the protected formats and media overwrites are not permitted, NetBackup takes the following actions:

If the volume has not been previously assigned for a backup	■ Sets the volume's state to FROZEN ■ Selects a different volume ■ Logs an error
If the volume is in the NetBackup media catalog and was previously selected for backups	■ Sets the volume's state to SUSPENDED ■ Aborts the requested backup ■ Logs an error

If the volume is mounted for a backup of the NetBackup catalog	The backup is aborted and an error is logged. The error indicates the volume cannot be overwritten.
If the volume is mounted to restore files or list the media contents	NetBackup aborts the request and logs an error. The error indicates that the volume does not have a NetBackup format.

Recommended use for Enable SCSI reserve property

All tape drive and bridge vendors support the SPC-2 SCSI reserve and release method. NetBackup has used SPC-2 SCSI reserve since NetBackup 3.4.3, and it is the default tape drive reservation method in NetBackup. SPC-2 SCSI reserve is effective for most NetBackup environments.

The SCSI persistent reserve method provides device status and correction and may be more effective in the following environments:

- Where NetBackup media servers operate in a cluster environment.
NetBackup can recover and use a reserved drive after a failover (if NetBackup owns the reservation). (With SPC-2 SCSI reserve, the drive must usually be reset because the reservation owner is inoperative.)
- Where the drive has high availability.
NetBackup can resolve NetBackup drive reservation conflicts and maintain high drive availability. (SPC-2 SCSI reserve provides no method for drive status detection.)

However, the SCSI persistent reserve method is not supported or not supported correctly by all device vendors. Therefore, thoroughly analyze the environment to ensure that all of the hardware supports SCSI persistent reserve correctly.

It is recommended to carefully consider all of the following factors before **Enable SCSI reserve** is used:

- Only a limited number of tape drive vendors support SCSI persistent reserve.
- SCSI persistent reserve is not supported or not supported correctly by all Fibre Channel bridge vendors. Incorrect support in a bridge means no access protection. Therefore, if the environment uses bridges, do not use SCSI persistent reserve.
- If parallel SCSI buses are used, carefully consider the use of SCSI persistent reserve. Usually, parallel drives are not shared, so SCSI persistent reserve protection is not required. Also, parallel drives are usually on a bridge, and bridges do not support SCSI persistent reserve correctly. Therefore, if the environment uses parallel SCSI buses, do not use SCSI persistent reserve.

- The operating system tape drivers may require extensive configuration to use SCSI persistent reserve. For example, if the tape drives do not support SPC-3 Compatible Reservation Handling (CRH), ensure that the operating system does not issue SPC-2 reserve and release commands.

If any of the hardware does not support SCSI persistent reserve, it is not recommended that SCSI persistent reserve is used.

Network properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the client. If necessary click **Connect**, then click **Edit client**. Click **Windows client > Network**.

Use the **Network** properties to configure the communications requirements between clients and the primary server. These properties apply to the currently selected Windows client .

The **Network** host properties contain the following settings.

Table 7-40 Network properties for Windows clients

Property	Description
NetBackup client service port (BPCD)	This property specifies the port that the NetBackup client uses to communicate with the NetBackup server. The default is 13782. Note: If you change this port number, remember that it must be the same for all NetBackup servers and clients that communicate with one another.
NetBackup request service port (BPRD)	This property specifies the port for the client to use when it sends requests to the NetBackup request service (bprd process) on the NetBackup server. The default is 13720. Note: If you change this port number, remember that it must be the same for all NetBackup servers and clients that communicate with one another.
Announce DHCP interval	This property specifies how many minutes the client waits before it announces that a different IP address is to be used. The announcement occurs only if the specified time period has elapsed and the address has changed since the last time the client announced it.

Network settings properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server or client. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Network settings**.

The **Network settings** host properties apply to primary servers, media servers, and clients.

The **Network settings** page contains properties for **Reverse host name lookup** and **Use the IP address family**.

See [“Reverse host name lookup property”](#) on page 163.

See [“Use the IP address family property”](#) on page 164.

Reverse host name lookup property

The domain name system (DNS) reverse host name lookup is used to determine what host and domain name a given IP address indicates.

Some administrators cannot or do not want to configure the DNS server for reverse host name lookup. For these environments, NetBackup offers the **Reverse host name lookup** property to allow, restrict, or prohibit reverse host name lookup.

Administrators can configure the **Reverse host name lookup** property for each host.

Table 7-41 **Reverse host name lookup property settings**

Property	Description
Allowed	The Allowed property indicates that the host requires reverse host name lookup to work to determine that the connection comes from a recognizable server. By default, the host resolves the IP address of the connecting server to a host name by performing a reverse lookup. If the conversion of the IP address to host name fails, the connection fails. Otherwise, it compares the host name to the list of known server host names. If the comparison fails, the host rejects the server and the connection fails.
Restricted	The Restricted property indicates that the NetBackup host first attempts to perform reverse host name lookup. If the NetBackup host successfully resolves the IP address of the connecting server to a host name (reverse lookup is successful), it compares the host name to the list of known server host names. If the resolution of the IP address to a host name fails (reverse lookup fails), based on the Restricted setting, the host converts the host names of the known server list to IP addresses (using a forward lookup). The host compares the IP address of the connecting server to the list of known server IP addresses. If the comparison fails, the host rejects the connection from server and the connection fails.

Table 7-41 Reverse host name lookup property settings (*continued*)

Property	Description
Prohibited	The Prohibited property indicates that the NetBackup host does not try reverse host name lookup at all. The host resolves the host names of the known server list to IP addresses using forward lookups. The NetBackup host then compares the IP address of the connecting server to the list of known server IP addresses. If the comparison fails, the NetBackup host rejects the connection from the server and the connection fails.

Use the IP address family property

On the hosts that use both IPv4 and IPv6 addresses, use the **Use the IP address family** property to indicate which address family to use:

- **IPv4 only** (Default)
- **IPv6 only**
- **Both IPv4 and IPv6**

While the **Use the IP address family** property controls how host names are resolved to IP addresses, the **Preferred network** properties control how NetBackup uses the addresses.

Nutanix AHV access hosts

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Nutanix AHV access hosts**.

You can also configure these settings in the web UI from **Workloads > Nutanix AHV**. Then select **AHV settings > Access hosts**.

Use the **Nutanix AHV access hosts** properties to configure a special host that is called a AHV access host. It is a NetBackup client that performs backups on behalf of the virtual machines.

For more information, see the [NetBackup for Nutanix AHV Administrator's Guide for details](#).

Port ranges properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server or client. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Port ranges**.

Use the **Port ranges** properties to determine how hosts connect to one another. These properties apply to the selected primary server, media server, or client.

The **Port ranges** host properties contain the following settings.

Table 7-42 Port ranges host properties

Property	Description
Use random port assignments	Specifies how the selected computer chooses a port when it communicates with NetBackup on other computers. Enable this property to let NetBackup randomly select ports from those that are free in the allowed range. For example, if the range is from 1023 through 5000, it chooses randomly from the numbers in this range. If this property is not enabled, NetBackup chooses numbers sequentially, not randomly. NetBackup starts with the highest number that is available in the allowed range. For example, if the range is from 1023 through 5000, NetBackup chooses 5000. If 5000 is in use, port 4999 is chosen. This property is enabled by default.
Client port window	Select Use OS selected non-reserved port to let the operating system determine which non-reserved port to use. Or, select the range of non-reserved ports on the selected computer. NetBackup can use any available port within this range as the source port when communicating with NetBackup on another computer.
Server port window	This property specifies the range of non-reserved ports on which NetBackup processes on this computer accept connections from NetBackup when the connection is not to a well known port. This property primarily applies to <code>bpcd</code> call-back when <code>vnetd</code> is disabled in the connect options and the local host name is configured for non-reserved ports. This property also applies in the situation where a third-party protocol is used, such as NDMP. It specifies the range of non-reserved ports on which this server accepts NetBackup connections from other computers. The default range is 1024 through 5000. Instead of indicating a range of ports, you can enable Use OS selected non-reserved port to let the operating system determine which non-reserved port to use. This setting applies to the selected primary or media server.

Table 7-42 Port ranges host properties (*continued*)

Property	Description
Server reserved port window	This entry specifies the range of local reserved ports on which this computer accepts connections from NetBackup when the connection is not to a well known port. This property primarily applies to <code>bpcd</code> call-back when <code>vnetd</code> is disabled in the connect options for a local host name. Instead of indicating a range of ports, you can enable Use OS selected non-reserved port to let the operating system determine which non-reserved port to use.

Registered ports and dynamically-allocated ports

NetBackup communicates between computers by using a combination of registered ports and dynamically-allocated ports.

Registered ports

These ports are registered with the Internet Assigned Numbers Authority (IANA) and are permanently assigned to specific NetBackup services. For example, the port for the NetBackup client daemon (`bpcd`) is 13782.

The following system configuration file can be used to override the default port numbers for each service:

On Windows: `%systemroot%\system32\drivers\etc\services`

On UNIX: `/etc/services`

Note: It is not recommended to change the port numbers that are associated with PBX (1556 and 1557).

Dynamically-allocated ports

These ports are assigned as needed, from configurable ranges in the **Port ranges** host properties for NetBackup servers and clients.

In addition to the range of numbers, you can specify whether NetBackup selects a port number at random or starts at the top of the range and uses the first one available.

Preferred network properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server or client. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Preferred network**.

Use the **Preferred network** properties to specify to NetBackup which networks or interfaces to use for outgoing NetBackup traffic from the selected hosts. These properties apply to currently selected primary server, media server, or client.

Note: The Preferred network setting in NetBackup does not apply to the Granular Recovery Technology (GRT) and VMware Instant Recovery features. Network settings that are configured in the operating system are used for these features during communication.

Preferred network entries are not needed if NetBackup is configured using host names with IP addresses to which the operating system resolves and then routes correctly.

When external constraints prevent the environment from being corrected, **Preferred network** entries may be useful as follows:

- Can be used to prevent NetBackup from connecting to specific destination addresses.
- Can be used to cause NetBackup to connect only to specific destination addresses.
- Can be used to request a subset of local interfaces for source binding when making outbound connections.

Caution: When used for source binding, the operating system may not honor the source binding list provided by NetBackup. If the operating system implements the weak host model, asymmetrical network routing may result. If asymmetrical routing occurs, the remote host may reject the inbound connection if it implements the strong host model. Similarly, stateful network devices may also drop asymmetrical connections. To ensure the use of specific outbound interfaces for specific remote hosts or networks, make sure that the OS name resolution and routing configurations are correct; create static host routes if needed. Ensure that all network drivers properly implement the IP and TCP networking protocols.

The local **Preferred network** entries do not affect the forwarding profile that the local host returns to a remote host during initial CORBA connection setup; it contains all the local plumbed interfaces. However, the End Point Selection algorithm within the remote process uses its local **Preferred network** entries to evaluate the profile when it selects the destination for the subsequent CORBA connection.

With respect to source binding, the **Preferred network** properties offer more flexibility than the **Use specified network interface** property in the **Universal settings** properties. The **Use specified network interface** property can be used to specify only a single interface for NetBackup to use for outbound calls. The

Preferred network properties were introduced so that administrators can give more elaborate and constrictive instructions that apply to multiple individual networks, or a range of networks. For example, an administrator can configure a host to use any network except one. If both properties are specified, **Use specified network interface** overrides **Preferred network**.

Note: Do not inadvertently configure hosts so that they cannot communicate with any other host. Use the `bptestnetconn` utility to determine whether the hosts can communicate as you intend.

See [“bptestnetconn utility to display Preferred network information”](#) on page 177.

The **Preferred network** host properties contain a list of networks and the directive that has been configured for each.

Table 7-43 Preferred network host properties

Property	Description
List of network specifications for NetBackup communications	The list of preferred networks contains the following information: ■ The Target column lists the networks (or host names or IP addresses) that have been given specific directives. If a network is not specifically listed as a target, or if a range of addresses does not include the target, NetBackup considers the target to be available for selection. Note that if the same network considerations apply for all of the hosts, the list of directives can be identical across all hosts in the NetBackup environment. If a directive contains an address that does not apply to a particular host, that host ignores it. For example, an IPv4-only host ignores IPv6 directives, and IPv6-only hosts ignore IPv4 directives. This action lets the administrator use the same Preferred network configurations for all the hosts in the NetBackup environment. ■ The Specified as column indicates the directive for the network: Match, Prohibited, or Only. ■ The Source column lists source binding information to use to filter addresses. The Source property is an optional configuration property.
Ordering arrows	Select a network in the list, then click the up or down arrow to change the order of the network in the list. The order can affect which network NetBackup selects. See “Order of directive processing in the Preferred network properties” on page 176.
Add	Click Add to add a network to the Preferred network properties. Then configure the directive for the network.
Actions > Edit	Locate a network in the list, then click Actions > Edit to change the Preferred network properties.

Table 7-43 Preferred network host properties (*continued*)

Property	Description
Actions > Delete	Locate a network in the list, then click Actions > Delete to remove the network from the list of preferred networks.

Add or edit a Preferred network setting

Refer to the following settings when you add or edit a preferred network setting.

Table 7-44 Configuration for Preferred network settings

Property	Description
Target	Enter a network address or a host name: ■ NetBackup recognizes the following wildcard entries as addresses: ■ 0.0.0.0 Matches any IPv4 address. ■ 0::0 Matches any IPv6 address. ■ 0/0 Matches the address of any family. ■ If the target is a host name which resolves to more than one IP address, only the first IP address will be used. ■ If a subnet is not specified, the default is /128 when the address is non-zero and /0 when the address is 0. This applies to both Target and Source properties. A subnet of /0 cannot be used with a non-zero address because it effectively negates all of the bits in the address, making the target or the source match every address. For example, 0/0. Note: Do not use the following malformed entries as wildcards: 0/32, 0/64, or 0/128. The left side of the slash must be a legitimate IP address. However, 0/0 may be used, as listed.
Match	The Match directive: ■ Applies when Target is a destination address. ■ Indicates that the specified network, address, or host name is preferred for communication with the selected host. ■ Does not reject other networks, addresses, or host names from being selected, even if they do not match. (The Only directive rejects unsuitable targets if they do not match.) ■ Is useful following a Prohibited or a Only directive. When used with other directives, Match indicates to NetBackup to stop rule processing because a suitable match has been found. ■ Can be used with the Source property to indicate source binding.

Table 7-44 Configuration for Preferred network settings (*continued*)

Property	Description
Prohibited	Use the Prohibited directive to exclude or prevent the specified network, address, or host name from being used. The Target is applied to both the source and the destination addresses. If a Source is specified and the Prohibited is indicated, it is ignored but the target is still prohibited. If the matched address is a destination address, evaluation stops. If this was the only potential destination, the connection is not attempted. If there are additional potential destinations, they are evaluated starting over with the first entry. If the matched address is a source address, it is removed from the source binding list. Caution: On some platforms, prohibiting a local interface may cause unexpected results when connecting to remote hosts. Prohibiting a local interface does not affect connections that are internal to the host.
Only	The Only directive: ■ Applies to destination addresses. ■ Indicates that the specified network, address, or host name that is used for communication with the selected host must be in the specified network. Use the Only directive to prevent any network from being considered other than those specified as Only. ■ If the address that is being evaluated does not match the target, it is not used and evaluation stops for that address. If the address being evaluated was the only potential destination, the connection is not attempted. If there is an additional potential destination, it is evaluated starting over with the first entry. ■ Can be used with the Source property to indicate source binding.
Source	Use this property with the Match or the Only directives to identify the local host name, IP addresses, or networks that may be used for source binding. If a subnet is not specified, the default is /128. If this host has an IP address that matches Source, that IP address will be used as the source when connecting to the destination. If the Source is not valid for this host, it is ignored.


How NetBackup uses the directives to determine which network to use

Each host has an internal table of preferred network rules that NetBackup consults before it selects a network interface to use for communication with another host. The table includes every interface-IP address combination available to the selected host. Based on the **Preferred NetBackup** directives, the table indicates to NetBackup whether or not the host is allowed to use a given network.

This topic uses the example of two multihomed servers (Server_A and Server_B) as shown in [Figure 7-1](#). Server A is considering which addresses it can use to access Server_B, given the **Preferred network** directives configured on Server_A.

When **Preferred network** directives are used to place restrictions on targets, they are added from the perspective of the server making the connection. The directives on Server_A affect its preferences as to which Server_B addresses it can use.

Figure 7-1 Multihomed servers example



[Figure 7-2](#) shows a table for Server_B. Server_B has multiple network interfaces, some of which have multiple IP addresses. In the table, yes indicates that NetBackup can use the network-IP combination as a source. In this example, no directives have been created for the host. Since no networks are listed in the **Preferred network** properties, any network-IP combinations can be used for communication.

Note: The following topic shows the `bptestnetconn` output for this example configuration:

See [“bptestnetconn utility to display Preferred network information”](#) on page 177.

Figure 7-2 From Server_A's perspective: Available IP addresses on Server_B when no directives are indicated on Server_A

		IP addresses	
		IPv4	IPv6
Network interfaces	2001:0db8:0:1f0::1efc	---	Yes
	10.80.73.147	Yes	---
	2001:0db8:0:11c::1efc	---	Yes
	2001:0db8:0:11d::1efc	---	Yes
	2001:0db8:0:11e::1efc	---	Yes
	10.96.73.253	Yes	---

Figure 7-3 shows a table for the same host (Server_B). Now, the **Preferred network** properties are configured so that all IPv4 addresses are excluded from selection consideration by NetBackup. All NetBackup traffic is to use only IPv6 addresses.

Figure 7-3 From Server_A's perspective: Available IP addresses on Server_B when directives to use IPv6 addresses only are indicated on Server_A

		IP addresses	
		IPv4	IPv6
Network interfaces	2001:0db8:0:1f0::1efc	---	Yes
	10.80.73.147	No	---
	2001:0db8:0:11c::1efc	---	Yes
	2001:0db8:0:11d::1efc	---	Yes
	2001:0db8:0:11e::1efc	---	Yes
	10.96.73.253	No	---

The following topics describe various configurations:

- See [“Configurations to use IPv6 networks”](#) on page 173.
- See [“Configurations to use IPv4 networks”](#) on page 175.
- See [“Configuration to prohibit using a specified address”](#) on page 178.
- See [“Configuration to prefer a specified address”](#) on page 179.
- See [“Configuration that restricts NetBackup to one set of addresses”](#) on page 180.

- See [“Configuration that limits the addresses, but allows any interfaces”](#) on page 181.

Configurations to use IPv6 networks

The following **Preferred network** configurations instruct NetBackup to use only IPv6 addresses as targets in outbound calls for the currently selected hosts. The configurations satisfy a topology where all backup traffic uses an IPv6 network and other traffic uses other networks.

One configuration uses the **Prohibited** directive ([Figure 7-4](#)) and one configuration uses the **Match** directive ([Figure 7-5](#)).

The more efficient method to specify one address family, (IPv6, in this case), is to prohibit IPv4. The behavior of the **Match** directive is not as exclusive as **Prohibited**. In this case, **Match** may not necessarily exclude other address families.

[Figure 7-4](#) uses the **Prohibited** directive with a wildcard to indicate to NetBackup to not consider using any IPv4 addresses. In this situation, NetBackup must use an IPv6 address.

Note: The default configuration is for NetBackup to use only IPv4 addresses.

If you have not previously changed the **Network settings > Use the IP address family** option to **Both IPv4 and IPv6** or **IPv6 only**, creating a directive that prohibits all IPv4 addresses renders the server mute.

See [“Use the IP address family property”](#) on page 164.

See [“Network settings properties”](#) on page 162.

Figure 7-4 Prohibit IPv4 addresses as targets

Add preferred network settings

Target
0.0.0.0

Specified as

☐ Match (The above network is preferred for communication)

☒ Prohibited (The above network is not used for communication)


☐ Only (Only target addresses in the above network is used for communication)

Source

Cancel Add and add another Add

Figure 7-5 uses the **Match** directive with a wildcard to indicate to NetBackup to prefer IPv6 addresses. In this case, NetBackup tries to use an IPv6 address, but may consider IPv4 addresses if necessary.

Figure 7-5 Match IPv6 addresses as targets



The screenshot shows a dialog box titled "Add preferred network settings". It has a close button (X) in the top right corner. The "Target" field contains the text "0::0". Below it, under the heading "Specified as", there are three radio buttons: "Match (The above network is preferred for communication)" which is selected, "Prohibited (The above network is not used for communication)", and "Only (Only target addresses in the above network is used for communication)". Below these is the "Source" field, which is currently empty. At the bottom of the dialog are three buttons: "Cancel", "Add and add another", and "Add".

Figure 7-6 shows another configuration that allows NetBackup to choose from multiple IPv6 networks.

Given the multihomed example configuration, the directive indicates the following:

- Four IPv6 networks, from `fec0:0:0:fe04` through `fec0:0:0:fe07`, are described as targets.
- For all addresses in these networks, a source binding address that is derived from the IP addresses of host name `host_fred` is used.

See [“How NetBackup uses the directives to determine which network to use”](#) on page 170.

Figure 7-6 Indicating a range of IPv6 networks

Add preferred network settings

Target
fec0:0:0:fe04::/62

Specified as

☐ Match (The above network is preferred for communication)

☐ Prohibited (The above network is not used for communication)

☒ Only (Only target addresses in the above network is used for communication)

Source
host_fred

Cancel Add and add another Add


Configurations to use IPv4 networks

The following **Preferred network** configurations instruct NetBackup to use only IPv4 addresses as targets in outbound calls for the currently selected hosts. The configurations satisfy a topology where all backup traffic uses an IPv4 network and other traffic uses other networks.

One configuration uses the **Prohibited** directive (Figure 7-7) and one configuration uses the **Match** directive (Figure 7-8).

The more efficient method to specify one address family, (IPv4, in this case), is to prohibit IPv6. The behavior of the **Match** directive is not as exclusive as **Prohibited**. In this case, **Match** may not necessarily exclude other address families.

Figure 7-7 uses the **Prohibited** directive with a wildcard to indicate to NetBackup to not consider using any IPv6 addresses. In this situation, NetBackup must use an IPv4 address.

Figure 7-7 Prohibit IPv6 addresses as targets

Add preferred network settings

Target
0::0

Specified as

☐ Match (The above network is preferred for communication)


☒ Prohibited (The above network is not used for communication)

☐ Only (Only target addresses in the above network is used for communication)

Source

Cancel Add and add another Add

Figure 7-8 uses the **Match** directive with a wildcard to indicate to NetBackup to prefer IPv4 addresses. In this case, NetBackup tries to use an IPv4 address, but may consider IPv6 addresses if necessary.

Figure 7-8 Match IPv4 addresses as targets

Add preferred network settings

Target
0.0.0.0

Specified as

☒ Match (The above network is preferred for communication)

☐ Prohibited (The above network is not used for communication)

☐ Only (Only target addresses in the above network is used for communication)

Source

Cancel Add and add another Add

Order of directive processing in the Preferred network properties

NetBackup sorts all directives into decreasing order by the **Target** subnet length so that the more specific network specifications, such as complete host names or IP addresses, match first. (For example, a **Target** with a /24 subnet is processed before a **Target** with a /16 subnet.) In this way, NetBackup can honor host-specific overrides.

If multiple directives have the same length subnet, NetBackup looks at the order in which the directives are listed.

Use the up or down arrows to the right of the list to change the order of the directives.

NetBackup processes each resolved destination address and each prospective source address relative to the directives. Directives that contain addresses that do not apply to either host are ignored.

bptestnetconn utility to display Preferred network information

The `bptestnetconn` utility is available to administrators to test and analyze host connections. Use the preferred network option (`--prefnet` or `-p`) to display information about the preferred network configuration, along with the forward lookup information of a host on the server list.

For example, `bptestnetconn -v6 -p -s -H host1` displays the directives in the order in which NetBackup processes them, which may not be the order in which they are configured.

- The `bptestnetconn` command is described in the [NetBackup Commands Reference Guide](#).
- The following article contains best practices for using `bptestnetconn` command:

[Figure 7-9](#) shows the `bptestnetconn` output when run on Server_A, for Server_B. That is, `bptestnetconn` is run from Server_A's perspective. Based on the directives configured on Server_A, for Server_B, `bptestnetconn` shows the available IP addresses on Server_B. In this example, no directives are configured on Server_A.

Figure 7-9 `bptestnetconn` for Server_B with no directives listed

```
[root@Server_A netbackup]# bptestnetconn -f --prefnet -H Server_B
```

FL: Server_B ->	10.81.73.147	:	11 ms SRC: ANY
FL: Server_B ->	10.96.73.253	:	11 ms SRC: ANY
FL: Server_B ->	2001:db8:0:11d::1efc	:	11 ms SRC: ANY
FL: Server_B ->	2001:db8:0:11e::1efc	:	11 ms SRC: ANY
FL: Server_B ->	2001:d8b:0:1f0::1efc	:	11 ms SRC: ANY
FL: Server_B ->	2001:db8:0:11c::1efc	:	11 ms SRC: ANY

```
-----
Total elapsed time: 0 sec
```

Host for which lookup
is performed

List of networks available to
Server_B

Any source is available to
use for a connection

The following directive is added to the **Preferred network** properties on Server_A:

In the configuration file the directive appears as follows:

```
PREFERRED_NETWORK = 2001:0db8:0:11c::/62 ONLY
```

This directive provides NetBackup with the information to filter the addresses and choose to communicate with only those that match the :11c, :11d, :11e, and :11f networks. The addresses that do not match the **Only** directive are prohibited, as shown in the `bptestnetconn` output.

Figure 7-10 shows the `bptestnetconn` output for Server_B, given this directive.


Figure 7-10 `bptestnetconn` for Server_B with directive

```
[root@Server_A netbackup]# bptestnetconn -f --prefnet -H Server_B
-----
FL: Server_B -> 10.81.73.147          :      11 ms TGT PROHIBITED
FL: Server_B -> 10.96.73.253         :      11 ms TGT PROHIBITED
FL: Server_B -> 2001:db8:0:11d::1efc  :      11 ms SRC: ANY
FL: Server_B -> 2001:db8:0:11e::1efc  :      11 ms SRC: ANY
FL: Server_B -> 2001:d8b:0:1f0::1efc  :      11 ms TGT PROHIBITED
FL: Server_B -> 2001:db8:0:11c::1efc  :      11 ms SRC: ANY
-----
Total elapsed time: 0 sec
```

List of networks available to Server_B
Directives make some targets unavailable to Server_B

Configuration to prohibit using a specified address

Figure 7-11 shows a configuration that prohibits NetBackup from using the specified address, or in this case, addresses.

Figure 7-11 Prohibited target example

Add preferred network settings

Target
192.168.100.0/24

Specified as

☐ Match (The above network is preferred for communication)

☒ Prohibited (The above network is not used for communication)

☐ Only (Only target addresses in the above network is used for communication)

Source

Cancel Add and add another Add


Configuration to prefer a specified address

[Figure 7-12](#) shows a configuration that makes NetBackup prefer to use one range of destination addresses over others that might be available.

Other available destination addresses will only be used if one of the following is true:

- No destination address exists that is in this range, or
- A **Match** is specified for those addresses using a larger subnet mask, or
- A **Match** is specified for those addresses with a same length subnet mask and the address is ordered above this directive.

A **Prohibited** directive can be used to prevent the use of an address within this range. The **Prohibited** directive would need either a longer subnet mask, or a subnet mask of equal length with the **Prohibited** directive ordered above the **Match** directive. Additional **Match** directives may be used to indicate the additional backup networks that are allowed.

Figure 7-12 Match network selection with the source

Add preferred network settings

Target
192.168.100.0/24

Specified as


- ☒ Match (The above network is preferred for communication)
- ☐ Prohibited (The above network is not used for communication)
- ☐ Only (Only target addresses in the above network is used for communication)

Source

Cancel Add and add another Add

Configuration that restricts NetBackup to one set of addresses

[Figure 7-13](#) configures NetBackup to use only the specified range of destination addresses, and the allowed source addresses must also be in the same range. The only exception is if other directives with larger subnets are present, or with equal-length subnets but ordered above this one.

Figure 7-13 Only network selection with the same source binding address

Add preferred network settings

Target
192.168.100.0/24

Specified as

- ☐ Match (The above network is preferred for communication)
- ☐ Prohibited (The above network is not used for communication)
- ☒ Only (Only target addresses in the above network is used for communication)

Source
192.168.100.0/24

Cancel Add and add another Add

A host with the **Only** directive configured considers only those target addresses in the 192.168.100.0 subnet. Additionally, source binding to the local interface must be done on the 192.168.100.0 subnet.

Configuration that limits the addresses, but allows any interfaces

Figure 7-14 shows a configuration that allows only the addresses that start with the specified prefix to be considered. No source binding is specified, so any interface may be used.

Figure 7-14 Limiting the addresses, without any source binding

The screenshot shows a dialog box titled "Add preferred network settings". It has a close button (X) in the top right corner. The "Target" field contains the text "fec0:0:1::/48". Below this, under the heading "Specified as", there are three radio buttons: "Match (The above network is preferred for communication)", "Prohibited (The above network is not used for communication)", and "Only (Only target addresses in the above network is used for communication)". The "Only" radio button is selected. Below the radio buttons is a "Source" field, which is currently empty. At the bottom of the dialog, there are three buttons: "Cancel", "Add and add another", and "Add".

Properties setting in host properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server or client. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Properties**.

The host property **Properties** includes the following information about the selected host.

Table 7-45 Properties information for a host

Property name	Description
Host	The NetBackup client name of the host.
Operating system	The operating system and OS version on which the host is installed.
OS type	The type of OS.
Host type	The type of host: Primary server, media server, or client.
IP address	The IP address of the host.

RHV access hosts properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **RHV access hosts**.

You can also configure these settings in the web UI from **Workloads > RHV**. Then select **RHV settings > Access hosts**.

Use the **RHV access hosts** properties to add or remove RHV backup hosts. These properties apply to the currently selected primary server .

For more information, see the [NetBackup Red Hat Virtualization Administrator's Guide](#).

Resilient network properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server or client. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Resilient network**.

For media servers and clients, the **Resilient network** properties are read only. When a job runs, the primary server updates the media server and the client with the current properties.

The **Resilient network** properties let you configure NetBackup to use resilient network connections for backups and restores. A resilient connection allows backup and restore traffic between a client and a NetBackup media server to function effectively in high-latency, low-bandwidth networks such as WANs. The data travels across a wide area network (WAN) to media servers in a central datacenter.

NetBackup monitors the socket connections between the remote client and the NetBackup media server. If possible, NetBackup re-establishes dropped connections and resynchronizes the data stream. NetBackup also overcomes latency issues to maintain an unbroken data stream. A resilient connection can survive network interruptions of up to 80 seconds. A resilient connection may survive interruptions longer than 80 seconds.

The NetBackup Remote Network Transport Service manages the connection between the computers. The Remote Network Transport Service runs on the primary server, the client, and the media server that processes the backup or restore job. If the connection is interrupted or fails, the services attempt to re-establish a connection and synchronize the data.

NetBackup protects only the network socket connections that the NetBackup Remote Network Transport Service (`nbrntd`) creates. Examples of the connections that are not supported are:

- Clients that back up their own data (deduplication clients and SAN clients)
- Granular Recovery Technology (GRT) for Exchange Server or SharePoint Server
- NetBackup `nbfsd` process.

NetBackup protects connections only after they are established. If NetBackup cannot create a connection because of network problems, there is nothing to protect.

Resilient connections apply between clients and NetBackup media servers, which includes primary servers when they function as media servers. Resilient connections do not apply to primary servers or media servers if they function as clients and back up data to a media server.

Resilient connections can apply to all of the clients or to a subset of clients.

Note: If a client is in a subdomain that is different from the server subdomain, add the fully qualified domain name of the server to the client's hosts file. For example, `india.veritas.org` is a different subdomain than `china.veritas.org`.

When a backup or restore job for a client starts, NetBackup searches the **Resilient network** list from top to bottom looking for the client. If NetBackup finds the client, NetBackup updates the resilient network setting of the client and the media server that runs the job. NetBackup then uses a resilient connection.

Table 7-46 Resilient network properties

Property	Description
FQDN or IP address	The full qualified domain name or IP address of the host. The address can also be a range of IP addresses so you can configure more than one client at once. You can mix IPv4 addresses and ranges with IPv6 addresses and subnets. If you specify the host by name, it is recommended that you use the fully qualified domain name. Use the arrow buttons on the right side of the pane to move up or move down an item in the list of resilient networks.
Resiliency	Resiliency is either On or Off .

Note: The order is significant for the items in the list of resilient networks. If a client is in the list more than once, the first match determines its resilient connection status. For example, suppose you add a client and specify the client IP address and specify **On** for **Resiliency**. Suppose also that you add a range of IP addresses as **Off**, and the client IP address is within that range. If the client IP address appears before the address range, the client connection is resilient. Conversely, if the IP range appears first, the client connection is not resilient.

Other NetBackup properties control the order in which NetBackup uses network addresses.

The NetBackup resilient connections use the SOCKS protocol version 5.

Resilient connection traffic is not encrypted. It is recommended that you encrypt your backups. For deduplication backups, use the deduplication-based encryption. For other backups, use policy-based encryption.

Resilient connections apply to backup connections. Therefore, no additional network ports or firewall ports must be opened.

Note: If multiple backup streams run concurrently, the Remote Network Transport Service writes a large amount of information to the log files. In such a scenario, it is recommended that you set the logging level for the Remote Network Transport Service to 2 or less. Instructions to configure unified logs are in a different guide.

View the resiliency status of a client

You can view the resiliency status of a client on the **Clients** tab of a policy or in the host properties for a client.

To view the resiliency status of a client in a policy

- 1 In the **NetBackup web UI**, open a policy.
- 2 Select the **Clients** tab.
- 3 The **Resiliency** column shows the status for each client in the policy.

To view the resiliency status of a client in host properties

- 1 In the **NetBackup web UI**, select **Host > Host properties**.
- 2 Select the client. If necessary, click **Connect**, then click **Edit client**.
- 3 Select **Resilient network**.

The **Resiliency** column shows the status for the client.

About Resilient jobs

The Resilient jobs feature lets the media server's job processes continue to run during a service disruption with the primary server. Backup metadata is cached to a user-defined location while the primary server processes are disrupted. Once the primary server re-establishes connections to the active media server processes, the cached data is transferred, and the backup proceeds.

To determine if a job is resilient, search the job details for the text, "job is resilient". If this text is present, the job is resilient.

The Resilient jobs feature is enabled by default. This feature is only available for some policy types. Please review the current requirements and limitations:

- The resiliency feature is either enabled or disabled. Backup jobs run as resilient jobs only when resiliency is enabled.
- Resilient jobs are only supported for Windows and Standard policy types.
- Backups cannot be multiplexed.
- Backups cannot have parent and child hierarchy. Use the Activity monitor to show parent and child relationship.
- Resilient jobs support the failure of the primary server. If the media server fails for any reason, the resilient jobs feature is not supported.

Note: If the primary server is also either the media server or the client, and it fails, the job is not resilient.

- If the client fails for any reason, the resilient job feature is not supported.
- If the primary server is upgraded while a backup is active, the backup is not resilient.
- The media server must be at NetBackup version 10.1.1 or later.
- Multistreamed backup jobs are not supported.
- Fiber Transport Media Server (FTMS) environments are not supported.

Resilient connection resource usage

Resilient connections consume more resources than regular connections, as follows:

- More socket connections are required per data stream. Three socket connections are required to accommodate the Remote Network Transport Service that runs on both the media server and the client. Only one socket connection is required for a non-resilient connection.

- More sockets are open on media servers and clients. Three open sockets are required rather than one for a non-resilient connection. The increased number of open sockets may cause issues on busy media servers.
- More processes run on media servers and clients. Usually, only one more process per host runs even if multiple connections exist.
- The processing that is required to maintain a resilient connection may reduce performance slightly.

Specify resilient connections for clients

Use the following procedure to specify resilient connections for NetBackup clients.

See [“Resilient network properties”](#) on page 182.

Alternatively, you can use the `resilient_clients` script to specify resilient connections for clients:

- Windows: `install_path\NetBackup\bin\admincmd\resilient_clients`
- UNIX: `/usr/opensv/netbackup/bin/admincmd/resilient_clients`

To specify resilient connections for clients

- 1 Open the **NetBackup web UI**.
- 2 On the left, click **Hosts > Host properties**.
- 3 Select the primary server. If necessary, click **Connect**. Then click **Edit primary server**.
- 4 Click **Resilient network**.
- 5 You can perform the following actions:

Add a setting

To add a host or IP address setting

- 1 Click **Add**.
- 2 Enter a client host name or an IP address.

If you specify the client host by name, it is recommended that you use the fully qualified domain name.
- 3 Ensure that the **On** option is selected.
- 4 Click **Add and add another**.
- 5 Repeat until you have added each setting.
- 6 When you finish adding network settings, click **Add**.

Edit a setting	To edit a host or IP address setting	
	1	Locate the client host name or the IP address.
	2	Click Actions > Edit .
	3	Select the desired Resiliency setting.
	4	Click Save .
Delete a setting	Delete a host or IP address setting	
	1	Locate the client host name or the IP address.
	2	Click Actions > Delete .
Up arrow, Down arrow	Change the order of items	
	1	Select the client host name or the IP address.
	2	Click the Up or Down button.
	The order of the items in the list is significant.	
	See "Resilient network properties" on page 182.	
The settings are propagated to the affected hosts through normal NetBackup inter-host communication, which can take up to 15 minutes.		
6	If you want to begin a backup immediately, restart the NetBackup services on the primary server.	

Resource limit properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Resource limits**.

The **Resource limits** properties control the number of simultaneous backups that can be performed on a particular resource type. These settings apply to all policies for the currently selected primary server.

Note: The **Resource limit** properties apply only to policies that use automatic selection of virtual machines (the policy's Query Builder). If you select virtual machines manually, the **Resource limit** properties have no effect.

See the respective guide for the workload or agent for details on the available resource limit properties.

Restore failover properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Restore failover**.

The **Restore failover** properties control how NetBackup performs automatic failover to a NetBackup media server. A failover server may be necessary if the regular media server is temporarily inaccessible to perform a restore operation. The automatic failover does not require administrator intervention. By default, NetBackup does not perform an automatic failover. These properties apply to currently selected primary servers.

The **Restore failover** host properties contain the following settings.

Table 7-47

Property	Description
Media server	Displays the NetBackup media servers that have failover protection for restores.
Failover restore servers	Displays the servers that provide the failover protection. NetBackup searches from top to bottom in the column until it finds another server that can perform the restore.

A NetBackup media server can appear only once in the **Media server** column but can be a failover server for multiple other media servers. The protected server and the failover server must both be in the same primary and media server cluster.

The following situations describe examples of when to use the restore failover capability:

- Two or more media servers share a robot and each has connected drives. When a restore is requested, one of the servers is temporarily inaccessible.
- Two or more media servers have standalone drives of the same type. When a restore is requested, one of the servers is temporarily inaccessible.

In these instances, inaccessible means that the connection between `bprd` on the primary server and `bptm` on the media server (through `bpcd`) fails.

Possible reasons for the failure are as follows:

- The media server is down.
- The media server is up but `bpcd` does not respond. (For example, if the connection is refused or access is denied.)

- The media server is up and `bpcd` is running, but `bptm` has problems. (For example, `bptm` cannot find the required tape.)

Assigning an alternate media server as a failover restore server

You can assign another media server to act as a failover restore server for your media server. If your media server is unavailable during a restore, the failover restore server takes its place.

To assign an alternate media server as a failover restore server

- 1 In the **NetBackup web UI** click **Hosts > Host properties**.
- 2 Select the primary server.
- 3 If necessary, click **Connect**. Then click **Edit primary server**.
- 4 Click **Restore failover**.
- 5 Click **Add**.
- 6 In the **Media server** field, specify the media server for failover protection.
- 7 In the **Failover restore servers** field, specify the media servers to try if the server that is designated in the **Media server** field is unavailable. Separate the names of multiple servers with a single space.
- 8 Click **Add**.
- 9 Click **Save**.

Before the change takes effect, you must stop and restart the NetBackup Request Daemon on the primary server where the configuration was changed.

Retention periods properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Retention periods**.

Use the **Retention periods** properties to define a duration for each retention level. You can select from 0-100 retention levels.

In a policy, the retention period determines how long NetBackup retains the backups or the archives that are created according to the schedule. These properties apply to selected primary servers.

By default, NetBackup stores each backup on a volume that already contains backups at the same retention level. However, NetBackup does not check the retention period that is defined for that level. When the retention period for a level

is redefined, some backups that share the same volume may have different retention periods.

For example, if the retention level 3 is changed from one month to 6 months, NetBackup stores future level 3 backups on the same volumes. That is, the backups are placed on the volumes with the level 3 backups that have a retention period of one month.

No problem exists if the new and the old retention periods are of similar values. However, before a major change is made to a retention period, suspend the volumes that were previously used for that retention level.

Note: If a backup or a duplicate job is configured with a retention level greater than 25 and a policy has a storage unit that is managed by a pre-NetBackup 8.0 media server, the backup jobs that are associated with the policy fail with the following error message:

```
Retention level <number> is not valid.
```

As a workaround, you can either upgrade the media server to NetBackup 8.0 or later or set the retention level between 0 and 25 in the policy. Note that the retention period for level 25 is always set to expire immediately and this value cannot be changed.

Note: For a manual import, if a primary or a media server that runs an earlier version than NetBackup 8.0 imports a backup image that was created on a NetBackup 8.0 primary server and configured with a retention level greater than 24, the import job resets the retention level to 9 (infinite). As a workaround, you can import such backup images from a primary or a media server that runs NetBackup 8.0 or later.

See [“Determining retention periods for volumes”](#) on page 192.

The **Retention periods** host properties contain the following settings.

Table 7-48 Retention periods page properties

Property	Description
Retention level	The retention level number (0 through 100).
	Value
	Assigns a number to the retention level setting.
	Units
	Specifies the units of time for the retention period. The list includes hours as the smallest unit of granularity and the special units, Infinite , and Expires immediately .

Table 7-48 Retention periods page properties (*continued*)

Property	Description
Retention period	A list of the current definitions for the possible levels of retention. By default, levels 9 through 100 (except level 25) are set to infinite. Retention level 9 cannot be changed and the retention period is always set to infinite. Retention level 25 also cannot be changed and the retention period is always set to expire immediately. See “Retention Periods with end dates beyond 2038, excluding Infinity” on page 192. With the default, there is no difference between a retention level of 12 and a retention level of 20, for example. If the retention period is changed for a level, it affects all schedules that use that level. The Changes pending column uses an asterisk (*) to indicate that the period has been changed and not applied. NetBackup does not change the actual configuration until the administrator accepts or applies the changes.
Schedule count	Lists the number of schedules that use the currently selected retention level.
Changes pending	This column displays an asterisk (*) to indicate that the period has been changed and not applied. NetBackup does not change the actual configuration until the administrator accepts or applies the changes.
Schedules using this retention level	Displays a list of the current policy names and schedule names that use the retention level.
Impact report	Displays a summary of how changes affect existing schedules. The list displays all schedules in which the retention period is shorter than the frequency period.

Changing a retention period

Use the following procedure to change a retention period.

To change a retention period

- 1 On the left, select **Hosts > Host properties**.
- 2 Select the primary server.
- 3 If necessary, click **Connect**. Then click **Actions > Edit primary server**.
- 4 Click **Retention periods**.

- 5 Locate the retention level to change and click **Edit**.

By default, levels 9 through 100 (except level 25) are set to infinite. If the levels are left at the default, there is no difference between a retention level of 12 and a retention level of 20. Level 9 cannot be changed and the retention period is always set to infinite. Retention level 25 also cannot be changed and the retention period is always set to expires immediately.

See [“Retention Periods with end dates beyond 2038, excluding Infinity”](#) on page 192.

The dialog box displays the names of all schedules that use the selected retention level as well as the policy to which each schedule belongs.

- 6 Type the new retention period in the **Value** box.
- 7 From the **Units** drop-down list, select a unit of measure (days, weeks, months, years, infinite, or expires immediately).

After you change the value or unit of measure, an asterisk (*) appears in the **Changes pending** column to indicate that the period was changed. NetBackup does not change the actual configuration until the administrator accepts or applies the changes.

- 8 Click **Impact report**.

The policy impact list displays the policies and the schedule names where the new retention period is less than the frequency period. To prevent a potential gap in backup coverage, redefine the retention period for the schedules or change the retention or frequency for the schedule.

Determining retention periods for volumes

Use the following procedure to determine retention periods for volumes.

To determine retention periods for volumes

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**.
- 3 Click the **Volumes** tab. Find the volume in the list and examine the value in the **Retention period** column.

To see all volumes that have the same retention period, click the **Retention period** column header to sort the volumes by retention period.

Retention Periods with end dates beyond 2038, excluding Infinity

For NetBackup versions before 9.0, there is a retention period limitation. Due to UNIX epoch time and the year 2038 problem, any expiration time that exceeds

January 19, 2038 is automatically set to expire on January 19, 2038. The images with such expiration times will expire in January 19, 2038 regardless of what the original intent of the retention levels was.

This issue does not apply to retention levels for which the retention period is set to **Infinity**. NetBackup never expires media with a retention set to **Infinity** unless instructed to do so by the NetBackup administrator.

Starting with NetBackup version 9.0, retention periods that extend beyond the year 2038 are supported. This retention period support is applicable not only to images but tape media as well.

Some backup images that are created with earlier versions may have expiration dates of January 19, 2038 after upgrade. You can correct the date issue with any of the images during upgrade or the records with end dates of January 19, 2038.

To correct the retention periods of infinity during upgrade, refer to the following article:

https://www.veritas.com/content/support/en_US/article.100048600

To correct the records with end dates of January 19, 2038, refer to the following article:

https://www.veritas.com/content/support/en_US/article.100048744

Scalable Storage properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the media server. If necessary click **Connect**, then click **Edit media server**. Click **Scalable storage**.

The **Scalable Storage** properties contain information about encryption, metering, bandwidth throttling, and network connections between the NetBackup hosts and your cloud storage provider. These properties appear only if the host is supported for cloud storage. See the *NetBackup Enterprise Server and Server - Hardware and Cloud Storage Compatibility List* for your release available through the following URL:

<http://www.netbackup.com/compatibility>

The **Scalable storage** properties apply to currently selected media server .

The **Scalable storage** host properties contain the following settings.

Table 7-49 Scalable storage host properties

Property	Description
Key Management Server (KMS) name	If you configured a key management service (KMS) server, the name of the primary server that sends the request to the KMS server is displayed here.
Metering interval	Determines how often NetBackup gathers connection information for reporting purposes. The value is set in seconds. The default setting is 300 seconds (5 minutes). If this value is set to zero, metering is disabled.
Total available bandwidth	Use this value to specify the speed of your connection to the cloud. The value is specified in kilobytes per second. The default value is 102400 KB/sec.
Sampling interval	The time, in seconds, between measurements of bandwidth usage. The larger this value, the less often NetBackup checks to determine the bandwidth in use. If this value is zero, throttling is disabled.
Advanced settings	Expand Advanced settings to configure additional settings for throttling. See “Configuring advanced bandwidth throttling settings” on page 194. See “Advanced bandwidth throttling settings” on page 195.
Maximum concurrent jobs	The default maximum number of concurrent jobs that the media server can run for the cloud storage server. This value applies to the media server, not to the cloud storage server. If you have more than one media server that can connect to the cloud storage server, each media server can have a different value. Therefore, to determine the total number of connections to the cloud storage server, add the values from each media server. If you configure NetBackup to allow more jobs than the number of connections, NetBackup fails any jobs that start after the number of maximum connections is reached. Jobs include both backup and restore jobs. You can configure job limits per backup policy and per storage unit. Note: NetBackup must account for many factors when it starts jobs: the number of concurrent jobs, the number of connections per media server, the number of media servers, and the job load-balancing logic. Therefore, NetBackup may not fail jobs exactly at the maximum number of connections. NetBackup may fail a job when the connection number is slightly less than the maximum, exactly the maximum, or slightly more than the maximum. A value of 100 is generally not needed.

Configuring advanced bandwidth throttling settings

Advanced bandwidth throttling settings let you control various aspects of the connection between the NetBackup hosts and your cloud storage provider.

See [“Scalable Storage properties”](#) on page 193.

To configure advanced bandwidth throttling settings

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host properties**.
- 3 Select the media server.
- 4 If necessary, click **Connect**. Then click **Edit media server**.
- 5 Click **Scalable storage**.
- 6 Expand **Advanced settings**.
- 7 Configure the settings and then click **Save**.

See [“Advanced bandwidth throttling settings”](#) on page 195.

Advanced bandwidth throttling settings

The following table describes the advanced bandwidth throttling settings.

Table 7-50 Advanced throttling configuration settings

Property	Description
Read bandwidth	Use this field to specify the percentage of total bandwidth that read operations can use. Specify a value between 0 and 100. If you enter an incorrect value, an error is generated. If there is insufficient bandwidth to transmit the specified amount of data within a few minutes, restore or replication failures may occur due to time-outs. Consider the total load of simultaneous jobs on multiple media servers when you calculate the required bandwidth. Default value: 100 Possible values: 0 to 100

Table 7-50 Advanced throttling configuration settings (*continued*)

Property	Description
Write bandwidth	Use this field to specify the percentage of total bandwidth that write operations can use. Specify a value between 0 and 100. If you enter an incorrect value, an error is generated. If there is insufficient bandwidth to transmit the specified amount of data within a few minutes, backup failures may occur due to time-outs. Consider the total load of simultaneous jobs on multiple media servers when you calculate the required bandwidth. Default value: 100 Possible values: 0 to 100
Work time	Use this field to specify the time interval that is considered work time for the cloud connection. Specify a start time and end time. Indicate how much bandwidth the cloud connection can use in the Allocated bandwidth field. This value determines how much of the available bandwidth is used for cloud operations in this time window. The value is expressed as a percentage or in kilobytes per second.
Off time	Use this field to specify the time interval that is considered off time for the cloud connection. Specify a start time and end time. Indicate how much bandwidth the cloud connection can use in the Allocated bandwidth field. This value determines how much of the available bandwidth is used for cloud operations in this time window. The value is expressed as a percentage or in kilobytes per second.
Weekend	Specify the start and stop time for the weekend. Indicate how much bandwidth the cloud connection can use in the Allocated bandwidth field. This value determines how much of the available bandwidth is used for cloud operations in this time window. The value is expressed as a percentage or in kilobytes per second.
Read Bandwidth (KB/s)	This field displays how much of the available bandwidth the cloud storage server transmits to a NetBackup media server during each restore job. The value is expressed in kilobytes per second.

Table 7-50 Advanced throttling configuration settings (*continued*)

Property	Description
Write Bandwidth (KB/s)	This field displays how much of the available bandwidth the NetBackup media server transmits to the cloud storage server during backup jobs. The value is expressed in kilobytes per second.

Servers properties

To access this setting, in the NetBackup web UI select **Hosts > Host properties**. Select the server or client. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Servers**.

The **Servers** properties display the NetBackup server lists on the selected primary server, media server, or client. The server lists display the NetBackup servers that the host recognizes.

The **Primary server** field contains the name of the primary server for the selected host. (The name of the selected host appears in the title bar.)

The **Servers** page contains the following settings.

Table 7-51 Servers properties

Tab	Description
Additional servers tab	This tab lists the additional servers that can access the server that is specified as the Primary server. During installation, NetBackup sets the primary server to the name of the system where the server software is installed. NetBackup uses the primary server value to validate server access to the client. The primary server value is also used to determine which server the client must connect to so that files can be listed and restored. Note: For a Fibre Transport (FT) media server that has multiple network interfaces for VLANs: Ensure that the FT server's primary host name appears before any other interface names for that FT media server host. For more information, see the NetBackup SAN Client and Fibre Transport Guide.
Media servers tab	This tab lists the hosts that are media servers only. Hosts that are listed as media servers can back up and restore clients, but have limited administrative privileges. If you add a media server to both the Media servers tab and the Additional servers tab, this action may introduce unintended consequences. A computer that is defined as both a primary server and a media server gives the administrator of the media server full primary server privileges. You may inadvertently give the media server administrator more privileges than intended.

Table 7-51 Servers properties (*continued*)

Tab	Description
Trusted primary servers tab	Use this tab to add the remote primary servers that you trust using NetBackup CA-signed certificates or external CA-signed certificates and to view the primary servers that are already trusted. See “Add a trusted primary server” on page 482. Note: If either the source or remote primary server is clustered, you must enable inter-node communication on all of the nodes in the cluster. Do so before you add the trusted primary server. See “Enable inter-node authentication for a NetBackup clustered primary server” on page 199. If your user account is configured for multifactor authentication on the target host, append the one-time password to the password.

Add a server to a servers list

Depending on the tab that is selected, you can add a primary server, media server, or client to the server list in the **Additional servers** tab or the **Media servers** tab.

To add a server to a servers list

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host properties**.
- 3 Select the host.
- 4 If necessary, click **Connect**. Then click **Edit primary server**, **Edit media server**, or **Edit client**.
- 5 Click **Servers**.
- 6 Select the tab that contains the server list that you want to modify.
- 7 Click **Add**.
- 8 Enter the name of the new server.
- 9 Click **Add**.

Note: If you add a media server, run `nbsmmcmd -addhost` to add the media server to the Enterprise Media Manager (EMM) in the NetBackup database of the primary server.

Remove a server from a servers list

You can remove a primary server or a media server from the **Additional servers** list or the **Media servers** list.

To remove a server from a servers list

- 1
- Open the NetBackup web UI.
- 2
- On the left, click **Hosts > Host properties**.
- 3
- Select the host.
- 4
- If necessary, click **Connect**. Then click **Edit primary server**, **Edit media server**, or **Edit client**.
- 5
- Click **Servers**.
- 6
- Click the **Additional servers** tab or the **Media servers** tab.
- 7
- Locate a server in the list.
- 8
- Click **Actions > Delete**.

Enable inter-node authentication for a NetBackup clustered primary server

NetBackup requires inter-node authentication among the primary servers in a cluster. For authentication, you must provision an authentication certificate on all of the nodes of the cluster. The certificates are used to establish SSL connections between the NetBackup hosts.

See [“Add a trusted primary server”](#) on page 482.

The inter-node authentication allows the following NetBackup functionality:

NetBackup web UI	The NetBackup web UI in primary server clusters requires the NetBackup authentication certificates for correct functionality.
Targeted A.I.R. (Auto Image Replication)	Auto Image Replication in which a primary server is in a cluster requires inter-node authentication among the hosts in that cluster. The NetBackup authentication certificates provide the means to establish the proper trust relationships. Provision the certificates on the cluster hosts before you add the trusted primary server. This requirement applies regardless of whether the clustered primary server is the source of the replication operation or the target.

To enable inter-node authentication for a NetBackup clustered primary server

- ◆ On the active node of the NetBackup primary server cluster, run the following NetBackup command:

- Windows: `install_path\NetBackup\bin\admincmd\bpnbaz -setupat`
- UNIX: `/usr/openv/netbackup/bin/admincmd/bpnbaz -setupat`

NetBackup creates the certificates on every node in the primary server cluster.

The following is example output:

```
# bpnbaz -setupat
You will have to restart Netbackup services on this machine after
the command completes successfully.
Do you want to continue(y/n)y
Gathering configuration information.
Please be patient as we wait for 10 sec for the security services
to start their operation.
Generating identity for host 'bitl.remote.example.com'
Setting up security on target host: bitl.remote.example.com
nbatd is successfully configured on Netbackup Primary Server.
Operation completed successfully.
```

Changing the primary server that performs backups and restores for a client

Use the **Make primary** option to change the primary server that performs backups and restores for a client. This option does not change a host into a primary server.

Note: The client can also change their primary server in the **Backup, Archive, and Restore** interface by selecting **Actions > Specify NetBackup Machines and Policy Type**. In this dialog, select the primary server to use for backups and restores.

This option is useful in a disaster recovery situation or in a NetBackup environment where Auto Image Replication is configured. For example, select a client in the source domain, then use the **Make primary** option to temporarily point the client to the primary server of the target domain. After you change the primary server, restores from the target domain can be initiated.

To change the primary server that a client uses for backups and restores

- 1 Open the NetBackup web UI.
- 2 On the left, click **Hosts > Host properties**.

- 3 Select the client.
- 4 If necessary, click **Connect**. Then click **Edit client**.
- 5 Click **Servers**.
- 6 On the **Additional servers** tab, locate the server.
- 7 Click **Actions > Make primary**.

In the configuration file, the new primary server appears as the first server entry in the list.

Changing the primary server does not prevent the former primary server from initiating backups for the client. As long as that server continues to be listed on the client's server list, the primary server can perform backups.

SharePoint properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the Windows client. If necessary click **Connect**, then click **Edit client**. Click **SharePoint**.

The **SharePoint** properties protect SharePoint Server installations and apply to the currently selected Windows client.

For complete information on these options, see the [NetBackup for Microsoft SharePoint Server Administrator's Guide](#).

The **SharePoint** host properties contain the following settings.

Table 7-52 SharePoint host properties

Property	Description
Domain\Username	Specifies the domain and the user name for the account you want to use to log on to SharePoint (DOMAIN\user name). Note: In 10.0 and later, credentials are stored in the Credential Management System (CMS).
Password	Specifies the password for the account.
Consistency check before backup	Specifies the consistency checks to perform on the SQL Server databases before NetBackup begins a backup operation. These checks are performed for both server-directed and user-directed backups. If you choose to perform a consistency check, you can select Continue with backup if consistency check fails . NetBackup then continues to perform the backup if the consistency check fails.

Table 7-52 SharePoint host properties (*continued*)

Property	Description
SharePoint granular restore proxy host	For any VMware backups that protect Federated SharePoint configurations, provide the name of the back-end SQL server. This server acts as the granular restore proxy host for the catalog hosts (front-end servers in the farm).

Consistency check options for SharePoint Server

The following consistency checks can be performed before a SharePoint Server backup.

Table 7-53 Consistency check options

Option	Description
None	Do not perform consistency checking.
Full check, excluding indexes	Select this option to exclude indexes from the consistency check. If indexes are not checked, the consistency check runs significantly faster but is not as thorough. Only the data pages and clustered index pages for each user table are included in the consistency check. The consistency of the non-clustered index pages is not checked.
Full check, including indexes	Include indexes in the consistency check. Any errors are logged.

SLP settings properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **SLP settings**. You can also configure the SLP settings from **Storage > Storage lifecycle policies > SLP settings**.

The **SLP settings** properties allow administrators to customize how storage lifecycle policies (SLPs) are maintained and how SLP jobs run. These properties apply to the SLPs of the currently selected primary server.

[Table 7-54](#) describes the available properties for SLPs. It also lists the syntax to use with the command-line method.

Use the list in the **Units** column to change the units of measurement for the size or the time.

Table 7-54 SLP settings

Property	Description
Minimum size per duplication job	The smallest batch size that can run as a single duplication job. The job does not run until enough images accumulate to reach this minimum batch size or until the Force interval for small jobs time is reached. Minimum: 1 kilobyte; no maximum size. Default: 8 gigabytes. Configuration option default: <code>SLP.MIN_SIZE_PER_DUPLICATION_JOB = 8 GB</code>
Maximum size per duplication job	The largest batch size that can run as a single duplication job. Minimum: 1 kilobyte; no maximum size. Default: 100 gigabytes. Configuration entry default: <code>SLP.MAX_SIZE_PER_DUPLICATION_JOB = 100 GB</code>
Maximum size per A.I.R. replication job	The largest batch size that can run as a single job for Auto Image Replication. Minimum: 1 kilobyte; no maximum size. Default: 100 gigabytes. Configuration entry default: <code>SLP.MAX_SIZE_PER_BACKUP_REPLICATION_JOB = 100 GB</code>
Maximum images per snapshot replication job	The largest number of images in a single batch that can run as a single job. Default: 50 images, with no minimum number or maximum number. Use this parameter with the Limit I/O streams disk pool option which limits the number of jobs that can run concurrently to each volume in the disk pool. Configuration entry default: <code>SLP.MAX_IMAGES_PER_SNAPSHOT_REPLICATION_JOB = 50</code>
Minimum images per A.I.R. Import job	The fewest number of images in a single batch that can run as an Auto Image Replication import job. The job does not run until either the minimum size is reached or the Force interval for small jobs time is reached. Minimum: 1 image; no maximum number of images. Default: 1 image. Configuration entry default: <code>SLP.MIN_IMAGES_PER_IMPORT_JOB = 1</code>
Maximum images per A.I.R. Import job	The largest number of images in a single batch that can run as an Auto Image Replication import job. Minimum: 1 job; no maximum number of images. Default: 250 images. Configuration entry default: <code>SLP.MAX_IMAGES_PER_IMPORT_JOB = 250</code>

Table 7-54 SLP settings (*continued*)

Property	Description
Force interval for small jobs	The age that the oldest image in a batch must reach after which the batch is submitted as a duplication job. This value prevents many small duplication jobs from running at one time or running too frequently. It also prevents NetBackup from waiting too long before it submits a small job. Default: 30 minutes, with no minimum number or maximum number. Configuration entry default: SLP.MAX_TIME_TIL_FORCE_SMALL_DUPLICATION_JOB = 30 MINUTES
Job submission interval	Indicates the frequency of the job submission for all operations. No minimum interval or maximum interval. Default: 5 minutes. By default, all jobs are processed before more jobs are submitted. Increase this interval to allow NetBackup to submit more jobs before all jobs are processed. Set the interval when the list of available images is scanned for those that can be batched together and jobs submitted. A shorter interval allows for a better response to changing system workloads at the cost of increased processing. Configuration entry default: SLP.JOB_SUBMISSION_INTERVAL = 5 MINUTES
Image processing interval	The number of minutes between image-processing sessions. Set the interval when newly created images are recognized and set up for SLP processing. Default: 5 minutes. Configuration entry default: SLP.IMAGE_PROCESSING_INTERVAL = 5 MINUTES
Cleanup interval	The time between when a job finishes and before NetBackup removes the job artifacts for the completed job. No minimum interval or maximum interval. Default: 24 hours. Configuration entry default: SLP.CLEANUP_SESSION_INTERVAL = 24 HOURS
Extended image retry interval	The amount of time to wait before an unsuccessful operation is added to the first job that runs after the delay. (This behavior applies to all SLP jobs.) The extra time gives the administrator additional time to solve a problem that prevents job completion. No minimum interval or maximum interval. Default: 2 hours. Configuration entry default: SLP.IMAGE_EXTENDED_RETRY_PERIOD = 2 HOURS

Table 7-54 SLP settings (*continued*)

Property	Description
Unused SLP definition version cleanup delay	Concerns the deletion of SLP versions where a more recent version exists. The setting controls how long a version must be inactive before NetBackup deletes it. Default: 14 days. Configuration entry default: <code>SLP.VERSION_CLEANUP_DELAY = 14 DAYS</code>
Tape resource multiplier	Limits the number of concurrently active duplication jobs that can access a single tape media storage unit to xx times the number of available drives. Allows tuning to avoid overloading the Resource Broker, yet makes sure that the devices are not idle. No minimum multiplier or maximum multiplier. Default: 2 (multiply access to the write drives by two). Configuration entry default: <code>SLP.TAPE_RESOURCE_MULTIPLIER = 2</code>
Disk resource multiplier	Limits the number of concurrently active duplication jobs that can access a single disk storage unit to xx times the number of available drives. Allows tuning to avoid overloading the Resource Broker, yet makes sure that the devices are not idle. No minimum multiplier or maximum multiplier. Default: 2 (multiply access to the write drives by two). Configuration entry default: <code>SLP.DISK_RESOURCE_MULTIPLIER = 2</code>
Group images across SLPs	If this parameter is set to Yes (default), multiple SLPs of the same priority can be processed in the same job. If No, batching can occur only within a single SLP. Configuration entry default: <code>SLP.DUPLICATION_GROUP_CRITERIA = 1</code> Configuration entry for no, do not allow batching: <code>SLP.DUPLICATION_GROUP_CRITERIA = 0</code>
Window close buffer time	Sets the amount of time before a window closes when NetBackup does not submit new jobs using that window. Minimum 2 minutes; maximum: 60 minutes. Default: 15 minutes. Configuration entry default: <code>SLP.WINDOW_CLOSE_BUFFER_TIME = 15 MINUTES</code>
Deferred duplication offset time	For deferred operations, jobs are submitted x time before the source copy is due to expire. Default: 4 hours. Configuration entry default: <code>SLP.DEFERRED_DUPLICATION_OFFSET_TIME = 4 HOURS</code>
Auto create A.I.R. Import SLP	Used for Auto Image Replication, indicates whether an SLP (that contains an Import operation) is created automatically in the target domain if no SLP is configured there. Default: Yes, an SLP is created in the target domain. Configuration entry default: <code>SLP.AUTO_CREATE_IMPORT_SLP = 1</code>

Table 7-54 SLP settings (*continued*)

Property	Description
How long to retry failed A.I.R. import jobs	How long NetBackup retries an Import job before it stops and deletes the record. After the initial four attempts, the retries become less frequent. Default: 0 (do not retry after the initial four attempts). Configuration entry default: <code>SLP.REPLICA_METADATA_CLEANUP_TIMER = 0 HOURS</code>
Pending A.I.R import threshold	How long NetBackup waits before it generates a notification that an Auto Image Replication copy is still in import pending state. After an Auto Image Replication copy has been replicated, NetBackup puts the source copy into import pending state. If the copy is in import pending state for the time period that this threshold sets, NetBackup generates a notification. Notifications are sent to the NetBackup error log and are visible in the Problems report. Notifications may also be sent to an email address, if specified. Default: 24 hours Configuration entry default: <code>SLP.PENDING_IMPORT_THRESHOLD = 24 HOURS</code>
Email address to receive notifications	The email address that receives pending A.I.R. import notifications. Default: None. Configuration entry format: <code>SLP.NOTIFICATIONS_ADDRESS = user@company.com</code>

Using the command line to change SLP parameters

You can also change the parameters using the command line.

To use the command-line method, use the `nbgetconfig` and the `nbsetconfig` commands to change the defaults. For information about these commands, see the [NetBackup Commands Reference Guide](#).

Command-line units of measurement for the SLP parameters

The abbreviations are case-insensitive for units of measurement.

The following abbreviations can be used where sizes are indicated:

bytes	kb	kilobyte	kilobyte(s)	kilobytes	mb	megabyte
megabyte(s)	megabytes	gb	gigabyte	gigabyte(s)	gigabytes	tb
terabyte	terabyte(s)	terabytes	pb	petabyte	petabyte(s)	petabytes

The following abbreviations can be used where units of time are indicated:

sec	second	second(s)	seconds	min	minute	minute(s)	minutes
hour	hour(s)	hours	day	day(s)	days	mon	month
month(s)	months	week	week(s)	weeks	year	year(s)	years

nbcl.conf file

Whenever a storage lifecycle policy parameter is changed from the default, the change creates the `nbcl.conf` configuration file.

This file is found in the following locations. It is present only if the default of any parameter has been changed.

- On Windows:
`install_path\NetBackup\var\global\nbcl.conf`
- On UNIX:
`/usr/opensv/var/global/nbcl.conf`

About batch creation logic in Storage Lifecycle Manager

The Storage Lifecycle Manager service (`nbstserv`) is in charge of creating duplication jobs for storage lifecycle policies. Part of duplication job creation includes grouping the backup (or source) jobs into batches.

Note: Restart `nbstserv` after making changes to the underlying storage for any operation in an SLP.

One objective of the batching logic is to prevent media contention for tape operations, including virtual tape libraries (VTL).

Batching logic applies to both disk and tape. (Though the method to prevent media contention for disk is to use disk pools and then to limit I/O streams to disk pools.)

The batching logic requires that for each evaluation cycle, `nbstserv` consider all completed source jobs when determining which duplication job to run next. By default, `nbstserv` performs the evaluation once every 5 minutes.

`nbstserv` avoids overloading the Resource Broker (`nbrb`) queue with jobs. Too many jobs in the queue make the role of the Resource Broker harder and slows down system performance.

By default, `nbstserv` now creates groups based on the **Group images across SLPs** parameter in the **SLP Parameters** host properties. By default, multiple storage lifecycle policies with the same priority can be batched together.

See [“SLP settings properties”](#) on page 202.

This batching logic change affects how duplication jobs appear in the **Activity Monitor**. Storage lifecycle policies that have been combined into one job appear under a single policy name: `SLP_MultipleLifecycles`. If a storage lifecycle policy has not been combined with another, the name appears in the **Activity Monitor** under the name of the SLP: `SLP_name`.

Users may see some duplication jobs that, although in the running state, do not duplicate data because they have no resources to read or write. These jobs continue to run until they receive resources to complete the job.

To turn off grouping by duplication job priority, set the **Group images across SLPs** parameter to **No** in the **SLP Parameters** host properties.

Throttle bandwidth properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **Throttle bandwidth**.

Use the **Throttle bandwidth** properties to specify a limit for the network bandwidth or transfer rate that NetBackup clients use on a network. The actual limiting occurs on the client side of the backup connection. These properties limit only backups. Restores are unaffected. The default is that the bandwidth is not limited.

The **Throttle bandwidth** properties are similar to the **Bandwidth** host properties, but offer greater flexibility in IPv6 environments.

To add, edit, or remove a throttle bandwidth setting

- 1 Open the **NetBackup web UI**.
- 2 On the left, click **Hosts > Host properties**.
- 3 Select the primary server. If necessary, click **Connect**. Then click **Edit primary server**.
- 4 Click **Throttle bandwidth**.

- Add a setting

To add a network or host setting

1

Click **Add**.

2

Enter the name of the network or host to which the throttle applies.

3

Select the bandwidth for the network or host indicated. A value of zero disables the throttling of IPv6 addresses.

This value is the transfer rate in kilobytes per second. A value of zero disables the throttling of IPv6 addresses.

4

Click **Add**.
- Edit a setting

To edit a network or host setting

1

Locate the name of the network or host.

2

Click **Actions > Edit**.

3

Make the wanted changes.

4

Click **Save**.
- Delete a setting

Delete a a network or host setting

1

Locate the name of the network or host.

2

Click **Actions > Delete**.

5 Click **Save**

Timeouts properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server or client. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Timeouts**.

The **Timeouts** properties apply to the selected primary server, media server, or client.

Table 7-55 Timeouts host properties

Property	Description
Client connect timeout	This property applies to the currently selected server. Specifies the number of seconds the server waits before it times out when it connects to a client. The default is 300 seconds.

Table 7-55 Timeouts host properties (*continued*)

Property	Description
Backup start notify timeout	This property applies to the currently selected server . Specifies the number of seconds the server waits for the <code>bpstart_notify</code> script on a client to complete. The default is 300 seconds. Note: If using the <code>bpstart_notify</code> script: The Client read timeout (<code>CLIENT_READ_TIMEOUT</code> option) must be equal to or greater than the Backup start notify timeout (<code>BPSTART_TIMEOUT</code> option). If the Client read timeout is less than the Backup start notify timeout, the job can time out while the <code>bpstart_notify</code> script is running.
Media server connect timeout	This property applies to the currently selected server . Specifies the number of seconds that the primary server waits before it times out when it connects to a remote media server. The default is 30 seconds.
Client read timeout	This property applies to the currently selected server or client. Specifies the number of seconds that NetBackup waits for a response from a client before the operation attempt fails. This timeout can apply to a NetBackup primary, remote media server, or database-extension client (such as NetBackup for Oracle). The default is 300 seconds. If the server does not get a response from a client within the Client read timeout period, the backup or the restore operation can fail. See the section called “Recommendations for the Client read timeout” on page 211. The sequence on a database-extension client is as follows: ■ NetBackup on the database-extension client reads the client’s client-read timeout to find the initial value. If the option is not set, the standard 5-minute default is used. ■ When the database-extension API receives the server’s value, it uses it as the client-read timeout.
Backup end notify timeout	This property applies to the currently selected server. Specifies the number of seconds that the server waits for the <code>bpend_notify</code> script on a client to complete. The default is 300 seconds. Note: If this timeout is changed, verify that Client read timeout is set to the same or higher value.

Table 7-55 Timeouts host properties (*continued*)

Property	Description
Use OS dependent timeouts	This property applies to the currently selected server or client. Specifies that the client waits for the timeout period as determined by the operating system when it lists files, as follows: ■ Windows client: 300 seconds ■ UNIX client: 1800 seconds File browse timeout Specifies how long the client can wait for a response from the NetBackup primary server while it lists files. If the limit is exceeded, the user receives a socket read failed error. The timeout can be exceeded even while the server processes the request. Note: If it exists, the value in a UNIX client's <code>\$HOME/bp.conf</code> file takes precedence to the property here.
Media mount timeout	This property applies to the currently selected primary server. Specifies how long NetBackup waits for the requested media to be mounted, positioned, and ready on backups, restores, and duplications. Use this timeout to eliminate excessive waiting time during manual media mounts. (For example, when robotic media is out of the robot or is off-site.)

Recommendations for the Client read timeout

It is recommended to increase the timeout value in the following situations:

- The client-read timeout on a database-extension client is a special case. Clients can initially require more time to get ready than other clients. More time is required because database backup utilities frequently start several backup jobs at the same time, slowing the central processing unit. A setting of 15 minutes is adequate for many installations.
- Backing up directly to an MSDP cloud storage server. If the value is not increased for both the primary server and the media server, you may see jobs failing with the following message in the job details:

```
Error bpbrm (pid=119850) socket read failed: errno = 62 - Timer
expired
```

Note that increasing the timeout is not needed if you use a storage lifecycle policy to first back up to an MSDP storage server and then duplicate the data to an MSDP cloud storage server using an optimized duplication operation. (This operation is the recommended method of operation.)

Note: If using the `bpstart_notify` script: The **Client read timeout** (`CLIENT_READ_TIMEOUT` option) must be equal to or greater than the **Backup start notify timeout** (`BPSTART_TIMEOUT` option). If the **Client read timeout** is less than the **Backup start notify timeout**, the job can timeout while the `bpstart_notify` script is running.

Universal settings properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the server or client. If necessary click **Connect**, then click **Edit primary server**, **Edit media server**, or **Edit client**. Click **Universal settings**.

Use the **Universal settings** properties to configure certain backup and restore settings. These properties apply to a selected primary server, media server, or client.

The **Universal settings** host properties contain the following settings.

Table 7-56 Universal settings properties

Property	Description
Restore retries	This setting applies to the selected server or client. Specifies the number of attempts a client has to restore after a failure. (The default is 0; the client does not attempt to retry a restore. The client can try up to three times.) Change Restore retries only if problems are encountered. If a job fails after the maximum number of retries, the job goes into an incomplete state. The job remains in the incomplete state as determined by the Move restore job from incomplete state to done state property. See “Clean up properties” on page 101. A checkpointed job is retried from the start of the last checkpointed file rather than at the beginning of the job. Checkpoint restart for restore jobs allows a NetBackup administrator to resume a failed restore job from the Activity Monitor.

Table 7-56 Universal settings properties (*continued*)

Property	Description
Browse timeframe for restores	This setting applies to the selected server and applies to all NetBackup clients. Specifies the timeframe that NetBackup uses to search for files to restore. By default, NetBackup includes files from the time of the last-full backup through the latest backup for the client. ■ Timeframe. Specifies how long ago NetBackup searches for files to restore. For example, to limit the browse range to one week before the current date, select Timeframe and specify 7. ■ Last full backup. Indicates whether NetBackup includes all backups since the last successful full backup in its browse range. This option is enabled by default. If the client belongs to more than one policy, then the browse starts with the earliest of the set of last-full backups.
Use specified network interface	This setting applies to the selected server or client. Specifies the network interface that NetBackup uses to connect to another NetBackup client or server. A NetBackup client or server can have more than one network interface. To force NetBackup connections to be made on a specific network interface, use this entry to specify the network host name of that interface. By default, the operating system determines the one to use.
Allow server file writes	This setting applies to the selected server or client. Specifies whether a NetBackup server can create or modify files on the NetBackup client. For example, disable this property to prevent server-directed restores and remote changes to the client properties. After the Allow server file writes property is applied, it can be cleared only by modifying the client configuration. The default is that server writes are allowed.
Administrator	This setting applies to the selected server or client. Specifies whether the server or the client sends email. ■ Server sends mail With this option the server sends an email to the address that is specified in the Global attributes properties. Enable this property if the client cannot send mail and you want an email notification. The default is that this property is disabled. See “Global attributes properties” on page 147. ■ Client sends mail With this option the client sends an email to the address that is specified in the Universal settings properties. If the client cannot send email, use Server sends mail. The default is that this property is enabled.

Table 7-56 Universal settings properties (*continued*)

Property	Description
Client administrator's email	Specifies the email address of the administrator on the client. This address is where NetBackup sends backup status reports for the client. By default, no email is sent. To enter multiple addresses or email aliases, separate entries with commas.

UNIX client properties

Use the **UNIX client** properties to define properties of clients running on the UNIX platform.

See [“Busy file settings properties”](#) on page 99.

See [“Client settings properties for UNIX clients”](#) on page 113.

See [“Lotus Notes properties”](#) on page 155.

UNIX Server properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the Linux primary server. If necessary click **Connect**, then and click **Edit primary server**. Then click **UNIX server**.

Use the **UNIX server** properties to change the **NFS access timeout** property. This property specifies how long the backup waits to process the mount table before it considers an NFS file system unavailable. The default is 5 seconds.

These properties apply to selected Linux primary servers.

User account settings properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **User account settings**.

Use the **User account settings** properties to customize the settings for user sessions, user account lockout, and the sign-in banner.

Table 7-57 User account settings properties

Property	Description
Session idle timeout	Logs out the user session if there is no activity for the specified period of time. See “Configure when idle sessions should time out” on page 464.

Table 7-57 User account settings properties (*continued*)

Property	Description
Maximum concurrent sessions	Limits the number of sessions that a user can have open concurrently. See “Configure the maximum of concurrent user sessions” on page 464.
User account logout	Lock out an account after the specified number of failed sign-in attempts. See “Configure the maximum of failed sign-in attempts” on page 464.
Sign-in banner configuration	You can configure a sign-in banner that displays each time that any user signs in to the NetBackup web UI. A different banner can be configured for any primary server. See “Display a banner to users when they sign in” on page 465.

VMware access hosts properties

To access this setting, in the web UI select **Hosts > Host properties**. Select the primary server. If necessary click **Connect**, then click **Edit primary server**. Click **VMware access hosts**.

You can also access this setting from **Workloads > VMware > VMware settings > Access hosts**.

Use the **VMware access hosts** host properties to add or remove VMware backup hosts. These properties apply to currently selected primary servers.

These properties appear when the NetBackup Enterprise Client license is installed.

The backup host is a NetBackup client that performs backups on behalf of the virtual machines. (This host was formerly known as the VMware backup proxy server.) The backup host is the only host on which NetBackup client software is installed. As an option, the backup host can also be configured as a NetBackup primary server or media server.

The backup host is referred to as the recovery host when it performs a restore

You can add servers to and remove servers from the access hosts list:

Add Click **Add** and enter the fully qualified domain name of the backup host.

Remove Locate the backup host in the list and click **Remove**.

For more information, see the [NetBackup for VMware Administrator's Guide](#).

Windows client properties

Use the **Windows client** properties to configure specific NetBackup properties for Windows clients.

See [“Client settings properties for Windows clients”](#) on page 117.

See [“Lotus Notes properties”](#) on page 155.

See [“Exchange properties”](#) on page 131.

See [“SharePoint properties”](#) on page 201.

See [“Active Directory properties”](#) on page 97.

See [“Enterprise Vault properties”](#) on page 129.

Configuration options not found in the host properties

Most NetBackup configuration options can be found in the **Host properties** of the **NetBackup web UI**. However, some options cannot be accessed in the **Host properties**.

To change the default value for an option that is not found in the **Host properties**, first use the `nbgetconfig` command to obtain a list of configuration options. Then use `nbsetconfig` to change the options as needed.

For information about these commands, see the following resources:

- [NetBackup Commands Reference Guide](#)
- See [“About using commands to change the configuration options on UNIX or Linux clients and servers”](#) on page 216.

For information about the configuration options that are available, see the [NetBackup Administrator's Guide, Volume I](#).

About using commands to change the configuration options on UNIX or Linux clients and servers

When commands (`nbsetconfig` or `bpsetconfig`) are used to change the configuration options on UNIX or Linux NetBackup servers or clients, the commands change the appropriate configuration files.

Most options are found in the following configuration file:


```
/usr/opensv/netbackup/bp.conf
```

If a single UNIX or Linux system is running as both a client and a server, the `bp.conf` file contains options for both the client and the server.

The `bp.conf` file observes the following syntax:

- Use the `#` symbol to comment out lines.
- Any number of spaces or tabs are allowed on either side of `=` signs.
- Blank lines are allowed.
- Any number of blanks or tabs are allowed at the start of a line.

Each nonroot user on a UNIX or Linux client can also have a personal `bp.conf` file in their home directory:

```
$HOME/bp.conf
```

The options in personal `bp.conf` files apply only to user operations. During a user operation, NetBackup checks the `$HOME/bp.conf` file before `/usr/opensv/netbackup/bp.conf`.

Root users do not have personal `bp.conf` files. NetBackup uses the `/usr/opensv/netbackup/bp.conf` file for root users.

Stop and restart all NetBackup daemons and utilities on the server after you make a change to the `bp.conf` file on a Linux primary server. This action ensures that all of the NetBackup processes use the new `bp.conf` values. This action is not required for changes to `bp.conf` files on a client or to a `$HOME/bp.conf` file on the primary server.

The `SERVER` option must be present in the `/usr/opensv/netbackup/bp.conf` file on all NetBackup UNIX or Linux clients and servers. During installation, NetBackup sets the `SERVER` option to the name of the primary server where the software is installed. It is the only required option in the `bp.conf` files. NetBackup uses internal software defaults for all options in the `bp.conf` file, except `SERVER`.

The `SERVER` entries must be the same on all servers in a primary and a media server cluster. It is recommended that all other entries also match on all servers. (The `CLIENT_NAME` option is an exception.)

Managing credentials for workloads and systems that NetBackup accesses

This chapter includes the following topics:

- [Overview of credential management in NetBackup](#)
- [Adding credentials in NetBackup](#)
- [Edit or delete a named credential](#)
- [Edit or delete Network Data Management Protocol \(NDMP\) credentials in NetBackup](#)
- [Add a configuration for an external CMS server](#)

Overview of credential management in NetBackup

Credential management lets you centrally manage the credentials that NetBackup uses to access systems and the workloads that it protects. You can manage NetBackup credentials, client credentials (for NDMP and disk arrays hosts), and External CMS server configurations from **Credential management**.

Credentials can be managed for the following workloads. To configure credentials for a workload (for example, SQL Server), refer to the guide for that workload for details.

Cassandra

MongoDB Ops Manager

Oracle

Cloud (for a cloud instance)

MSDP-C

PaaS database

Cloud object store	MySQL Server	PostgreSQL Server
Kubernetes	Nutanix AHV	SaaS
Microsoft SQL Server	Nutanix AHV Prism Central	

Credentials can also be managed for the following systems:

A Callhome proxy server	MSDP Samba User	Remote Primary Server
CyberArk	Malware detection (Malware scan host)	VMware guest VM
Disk arrays	Microsoft Sentinel	Veritas Alta Recovery Vault Azure
External Key Management Services (KMS)	NDMP	

Adding credentials in NetBackup

You can use the **Credential management** node to add a credential that NetBackup uses to connect to a system or workload.

- See [“Add a credential for NetBackup Callhome Proxy”](#) on page 219.
- See [“Add a credential for an external KMS”](#) on page 220.
- See [“Add a credential for Network Data Management Protocol \(NDMP\)”](#) on page 221.
- See [“Add a configuration for an external CMS server”](#) on page 223.

For SQL Server, Cloud, Kubernetes, and other workloads, refer to the guide for that workload for details.

[NetBackup documentation portal](#)

Add a credential for NetBackup Callhome Proxy

This type of credential provides the proxy server configuration that both the NetBackup Product Improvement Program and Usage Insights use.

See the [Cohesity Usage Insights Getting Started Guide](#) for details on using a Call Home proxy server.

To add a credential for NetBackup Callhome Proxy

- 1 On the left, click **Credential management**.
- 2 On the **Named credentials** tab, click **Add** and provide the following properties:
 - Credential name
 - Tag
 - Description
- 3 Click **Next**.
- 4 Select **Callhome proxy**.
- 5 Provide the credential details that are needed for authentication and click **Next**.
- 6 Add a role that you want to have access to the credential.
 - Click **Add**.
 - Select the role.
 - Select the credential permissions that you want the role to have.
- 7 Click **Next** and follow the prompts to complete the wizard.
- 8 After you create the credential, you must update the NetBackup configuration with an entry for `CALLHOME_PROXY_NAME`. Set the `CALLHOME_PROXY_NAME` to the credential name. From the primary server, use the command shown:

```
echo CALLHOME_PROXY_NAME = CredentialName |bpsetconfig.exe
```

Add a credential for an external KMS

This type of credential allows you to access an external KMS server that you have configured.

To add a credential for an external KMS

- 1 On the left, click **Credential management**.
- 2 On the **Named credentials** tab, click **Add** and provide the following properties:
 - Credential name
 - Tag
 - Description (for example: This credential is used to access the external KMS.)
- 3 Click **Next**.
- 4 Select **External KMS**.

5 Provide the credential details that are needed for authentication.

These details are used to authenticate the communication between the NetBackup primary server and the external KMS server:

- Certificate - Specify the certificate file contents.
- Private key - Specify the private key file contents.
- CA Certificate - Specify the CA certificate file contents.
- Passphrase - Enter the passphrase of the private key file.
- CRL check level - Select the revocation check level for the external KMS server certificate.

CHAIN - The revocation status of all the certificates from the certificate chain are validated against the CRL.

DISABLE - Revocation check is disabled. The revocation status of the certificate is not validated against the CRL during host communication.

LEAF - The revocation status of the leaf certificate is validated against the CRL.

See the [NetBackup Security and Encryption Guide](#) for more information on external KMS configuration.

6 Click **Next**.

7 Add a role that you want to have access to the credential.

- Click **Add**.
- Select the role.
- Select the credential permissions that you want the role to have.

8 Click **Next** and follow the prompts to complete the wizard.

Add a credential for Network Data Management Protocol (NDMP)

You can add the credentials that NetBackup uses to connect to the Network Data Management Protocol (NDMP).

For more information about NDMP credentials, see the [NetBackup NAS Administrator's Guide](#).

To add an NDMP credential

- 1** On the left, click **Credential management**.
- 2** Click the **Client credentials** tab.
- 3** Click **Add**.
- 4** Select **NDMP host** and click **Next**.

- 5 Enter an NDMP host name.
- 6 Select the type of host credential.
 - **Use the following credentials for this NDMP host on all media servers**
– This option uses the same credentials for all media servers.
 - **Use different credentials for this NDMP host on each media server** –
This option lets you enter unique credentials for each media server. After you enter credentials for each of the media servers, click **Add**.
- 7 Click **Add**.

Edit or delete a named credential

You can edit the properties for a named credential or delete a named credential NetBackup from the **Credential management**.

Edit a named credential

You can edit a named credential to change the following: credential tag, description, category, authentication details, or permissions. You cannot change the credential name.

To edit a named credential

- 1 On the left, select **Credential management**.
- 2 On the **Named credentials** tab, locate and select the check box for the credential that you want to edit.
- 3 Select **Edit** and update the credential as needed.
- 4 Review the changes and select **Finish**.
- 5 (Conditional) For any cloud workloads that use an agentless connection for instances, after you edit the credentials select the **Connect** button to reconnect the instances.

Delete a named credential

You can delete a named credential that you no longer need to use with NetBackup. Be sure to apply another credential to any assets that use the credential you want to delete. Otherwise, backups and restores may fail for those assets.

To delete a named credential

- 1 On the left, select **Credential management**.
- 2 On the **Named credentials** tab, locate and select the check box for the credential that you want to delete.

- 3 Select **Delete > Delete**.
- 4 (Conditional) If the credential deleted was a proxy credential, you must remove the `CALLHOME_PROXY_NAME` entity. From the primary server, use the following command to remove the `CALLHOME_PROXY_NAME` entity.

```
echo CALLHOME_PROXY_NAME |bpsetconfig.exe
```

Edit or delete Network Data Management Protocol (NDMP) credentials in NetBackup

You can edit or delete credentials for any media servers that use the Network Data Management Protocol (NDMP).

For more information about NDMP credentials, see the [NetBackup NAS Administrator's Guide](#).

Edit an NDMP credential

To edit an NDMP credential

- 1 On the left, click **Credential management**.
- 2 Click the **Client credentials** tab.
- 3 Locate and select the host. Then click **Edit**.
- 4 Make any changes that you want, then click **Save**.

Delete an NDMP credential

To delete an NDMP credential

- 1 On the left, click **Credential management**.
- 2 Click the **Client credentials** tab.
- 3 Select one or more hosts. Then click **Delete > Delete**.

Add a configuration for an external CMS server

This section provides you the procedure for adding a configuration for an external CMS server.

To add a configuration for an external CMS server

- 1 On the left, click **Credential management**.
- 2 On the **External CMS servers** tab, click **Add** and provide the following properties:

- Configuration name
- Description (for example: This configuration is used to access the external CMS.)
- External CMS provider
- Host name
- Port number: Default port number 443 would be considered (if not provided by the user).

Note: While configuring the external CMS server for CyberArk server, user can use the DNS hostname or IPV4 address. However it is recommended to use the DNS hostname for connecting to the host. CyberArk configuration fails if IPV6 address is used.

- 3 Click **Next**.
- 4 On the Associate credentials page, **Select existing credential** or **Add a new credential**.

More information is available on how to add a new credential.
See [“Add a credential for CyberArk”](#) on page 225.
- 5 Click **Next** and follow the prompts to complete the wizard.

Configure external credentials

This type of credential allows you to configure an external CMS server.

An **External** credential can only be created if an external CMS server configuration exists.

To configure external credentials

- 1 On the left, click **Credential management**.
- 2 On the **Named credentials** tab, click **Add**.
- 3 Select **External** and click **Start**.

On **Add a credential** page, provide the following properties:
 - Credential name
 - Tag
 - Description
- 4 Select the appropriate **Category** to assign the credential.

5 Search and select the **External CMS configuration**.

Provide the following parameter details for CyberArk Server:

- Application ID - Unique ID of the application issuing the password request.
- Object - Name of the password object to retrieve.
- Safe - Name of the Safe where the password is stored.

For more information on the parameters for CyberArk server, see [Call the Web Service using REST](#).

6 Click **Next**.

7 Add a role that you want to have access to the credential.

- Click **Add**.
- Select the role.
- Select the credential permissions that you want the role to have.

8 Click **Next** and follow the prompts to complete the wizard.

Add a credential for CyberArk

This type of credential allows you to access an external CMS server.

To add a credential for an external CMS server

1 On the left, click **Credential management**.

2 On the **Named credentials** tab, click **Add**.

3 Select **NetBackup** and click **Start**.

On **Add a credential** page, provide the following properties:

- Credential name
- Tag
- Description (for example: This credential is used to access the external CMS.)

4 Click **Next**.

5 Select **CyberArk** as the category.

6 Provide the credential details for CyberArk server:

These details are used to authenticate the communication between the NetBackup primary server and the external CMS server:

- Certificate - Specify the certificate file contents.

- Private key - Specify the private key file contents.
- CA Certificate - Specify the CA certificate file contents.
- Passphrase - Enter the passphrase of the private key file.
- CRL check level - Select the revocation check level for the external CMS server certificate.
 - CHAIN - The revocation status of all the certificates from the certificate chain are validated against the CRL.
 - DISABLE - Revocation check is disabled. The revocation status of the certificate is not validated against the CRL during host communication.
 - LEAF - The revocation status of the leaf certificate is validated against the CRL.

7 Click **Next**.

8 Add a role that you want to have access to the credential.

- Click **Add**.
- Select the role.
- Select the credential permissions that you want the role to have.

9 Click **Next** and follow the prompts to complete the wizard.

Certificate revocation lists for CyberArk server

Certificate revocation list (CRL) for an external certificate authority (CA) contains a list of digital certificates that the external CA has revoked before the scheduled expiration date and should no longer be trusted. NetBackup supports PEM and DER formats for CRLs for external CA. CRL's for all CRL issuers or external CA's are stored in the NetBackup CRL cache that resides on each host. During secure communication, NetBackup host verifies the revocation status of the peer host's external certificate with the CRL that is available in the NetBackup CRL cache, based on the CRL check level configuration option. For external CMS server, NetBackup supports CDP based server certificates.

NetBackup downloads the CRLs from the URLs that are specified in the peer host certificate's CDP and caches them in the NetBackup CRL cache.

To use CRL's from CDP:

- Ensure that the host can access the URLs that are specified in the peer host's CDP.
- Ensure that the **CRL check level** configuration option is set to a value other than **DISABLE**.

By default, CRLs are downloaded from the CDP after every 24 hours and updated in the CRL cache. To change the time interval, set the

ECA_CRL_REFRESH_HOURS configuration option to a different value. To manually delete the CRL's from the CRL cache, run the `nbcertcmd -cleanupCRLCache` command. The NetBackup CRL cache contains only the latest copy of a CRL for each CA (including root and intermediate CAs). The `bpcintcmd -crl_download` service updates the CRL cache during host communication in the following scenarios irrespective of the time interval set for the **ECA_CRL_REFRESH_HOURS** options:

- When CRLs in the CRL cache are expired.
- If CRLs are available in the CRL source, but they are missing from the CRL cache.

For details of **ECA_CRL_REFRESH_HOURS**, refer to **ECA_CRL_REFRESH_HOURS** for NetBackup servers and clients section from *NetBackup Security and Encryption Guide*.

Note: By default, the **ECMS_HOSTS_SECURE_CONNECT_ENABLED** flag is enabled (set to true). If this flag is enabled, the certificate deployed on the external CMS server must have Common Name or Subject Alternative Name that matches the host name of the external CMS server. Else, the connection to the external CMS server fails. For more information, see the **ECMS_HOSTS_SECURE_CONNECT_ENABLED** section in *NetBackup™ Administrator's Guide, Volume I*.

Edit or delete the configuration for an external CMS server

You can edit the properties of the configuration or delete the configuration from the **Credential management**.

Edit the configuration

You can edit the configuration to change the Description only. You cannot change the following properties: Configuration name, External CMS provider, Host name and Port number.

To edit the configuration

- 1 On the left, click **Credential management**.
- 2 On the **External CMS servers** tab, locate and click on the configuration that you want to edit.
- 3 Click **Edit** and update the properties as needed.
- 4 Review the changes and click **Next**.

- 5** Select an existing credential or add a new credential and click **Next**.
- 6** Review the changes and click **Finish**.

Delete the configuration

You can delete the configuration that you no longer need to use with NetBackup.

To delete the configuration

- 1** On the left, click **Credential management**.
- 2** On the **External CMS servers** tab, locate and click on the configuration that you want to delete.
- 3** Click **Delete**.
- 4** Confirm the deletion by clicking on **Remove**.

Troubleshooting the external CMS server issue

If the CyberArk Application ID contains internationalized characters and CyberArk server does not have the appropriate language pack installed, then the NetBackup user encounters a failure in adding the workload credentials from CyberArk.

Recommended action:

If the CyberArk Application Id contains internationalized characters, then install the corresponding language pack on the CyberArk server.

Managing deployment

This chapter includes the following topics:

- [About the deployment policies utility](#)
- [Managing the NetBackup Package repository](#)
- [Update host](#)
- [Deployment policies](#)
- [Copy a deployment policy](#)
- [Manually deploy a deployment policy](#)
- [Deployment job status](#)

About the deployment policies utility

Deployment policies are the main component of VxUpdate that serves as a client or host upgrade tool. The deployment policy lets you configure and run deployment activities on a schedule or enable the host owners to upgrade at their convenience. You can schedule precheck, staging, and installation tasks as separate activities with different schedules, each with their own specific deployment windows.

For more information regarding VxUpdate, see the *About VxUpdate* section within the [NetBackup Upgrade Guide](#).

Deployment policies are located in the NetBackup web UI under **Hosts > Deployment Management**. These policies provide the instructions that NetBackup follows to upgrade clients or hosts. Use this utility to provide the following instructions for a client or host upgrade:

What type of client or host to upgrade

See [“Attributes tab in Deployment management”](#) on page 232.

Which clients or hosts to upgrade

When to perform VxUpdate

See “[Schedules tab in Deployment management](#)” on page 234.

The security options to use for the clients or hosts

Managing the NetBackup Package repository

The NetBackup Package repository provides a central location to add and remove NetBackup packages. Packages let you upgrade NetBackup or deploy emergency engineering binaries in your NetBackup environment.

The interface arranges packages by NetBackup version number. For a specific version of NetBackup, there are multiple child packages, one for each supported platform.

Select **Hosts > Deployment management** to review the packages that are available to deploy to computers in your NetBackup environment. Actions available from this interface include:

- Add new packages.
- Delete existing packages.

Before you can add packages to the repository, you must download VxUpdate formatted packages from the myveritas.com licensing portal. Place downloaded package in an accessible location on the primary server. For details on how to download packages, see the **Repository management** section of *NetBackup Upgrade Guide*. Specifically, refer to the **Downloading Veritas NetBackup approved media server and client packages** procedure.

To add packages

- 1 From **Hosts > Deployment management**, select **Add package** or **Add**, depending if there are already packages in the repository.
- 2 Navigate to where your VxUpdate packages are located and select them. Be aware that NetBackup can only add the packages that reside on the primary server's file system.

The interface displays only VxUpdate packages. A directory may have files but if there are no VxUpdate packages, it shows as empty.

- 3 Select **Add** to add the packages.

Depending on the number and the size of packages you add, it may take a while for them to display in the repository.

To delete packages

- 1 From **Hosts > Deployment management**, select the packages you want to delete.
- 2 Select **Delete**.

Note: If you delete a parent package, all child packages that are associated with that parent are removed.

If you delete a server package, the associated client package is also deleted. For example, if you delete the Windows 8.3 server package, the Windows 8.3 client package is also removed.

Update host

The **Update host** option lets you launch immediate jobs to update or upgrade your NetBackup environment.

After you select **Hosts > Host properties** and make one or more valid selections, the **Update host** option appears in the upper right. Certain restrictions apply to the use of the **Update host** option:

- All computers you select must be of the same type. Select either all client computers or all media servers. If you select mixed computer types, the **Update host** option disappears.
- Primary servers are not supported. If you select a primary server, the **Update host** option disappears.
- The operating system and versions column must contain data for the **Update host** option to appear. If these columns do not contain data, attempt to connect to the host.

After you specify computers to update, select **Update host** to launch the update process. You are prompted for the information shown:

- **Attributes**

On this screen, specify: The package you want deployed, the operation type, any limit on concurrent jobs, and how to handle Java and the JRE.

See [“Attributes tab in Deployment management”](#) on page 232.

- **Hosts**

Displays the hosts you want to upgrade. From this screen, you can remove hosts.

See [“Hosts tab in Deployment management”](#) on page 233.

- **Security options** (if it appears)

Either accept the default (**Use existing certificates when possible**) or specify the appropriate security information for your environment.

See [“Security options tab in Deployment management”](#) on page 235.

- **Review**

Displays all the options you selected on previous screens.

Select **Update** to start the deployment job.

Deployment policies

Under **Hosts > Deployment management**, there is a **Deployment policies** tab. Use this tab to add, edit, copy, deactivate, delete, and launch your policies.

To add a new policy

- 1 Go to **Hosts > Deployment management > Deployment policies** and select **Add**.
- 2 Enter the required information for deployment policies.
See [“Attributes tab in Deployment management”](#) on page 232.
See [“Hosts tab in Deployment management”](#) on page 233.
See [“Schedules tab in Deployment management”](#) on page 234.
See [“Security options tab in Deployment management”](#) on page 235.
- 3 Select **Save**.

Similarly, to edit, copy, deactivate, or delete deployment policies, select the policy. Then select the appropriate action from banner.

To manually initiate policies, select the desired policy and select **Deploy now** from the menu.

Attributes tab in Deployment management

Use the policy **Attributes** tab to configure deployment management settings when you add a new deployment policy or change an existing deployment policy.

Setting	Description
Package	Select the package that you want to deploy. Note: You must add packages to the VxUpdate repository before you can create a working deployment policy. You can create deployment policies without packages in the repository, but those policies fail to run successfully. For more information regarding adding packages, see the <i>Repository Management</i> section within the NetBackup Upgrade Guide.
Media server	Specify the media server. This media server is used to connect and transfer files to the NetBackup hosts that are included in the policy. The media server must be version NetBackup 8.1.2 or later. Since the repository resides on the primary server, the primary server is the default value for the media server field.
Limit simultaneous jobs	Select the Limit simultaneous jobs option and specify a value for jobs to limit the total number of concurrent jobs that can run at one time. The default value is 3. The minimum value is 1 and the maximum value is 999. If you want to set unlimited simultaneous upgrade jobs, you must specify a value which is equivalent or higher than the count of the number of hosts that are selected for upgrade. For example, if you have selected 50 hosts, ensure that the Limit simultaneous jobs value is set to 50 or more but lower than the maximum value which is 999.
Java GUI and JRE	Specify if you want the NetBackup Administration Console and the JRE upgraded on the target systems. The three options include: ■ Match: Preserve the current state of the NetBackup Administration Console and JRE components. The components are upgraded if they are present on the pre-upgraded system. The components are not installed if they are not present on the pre-upgraded system. ■ Include: Install or upgrade the NetBackup Administration Console and JRE components on the specified computers. ■ Exclude: Exclude the NetBackup Administration Console and JRE components from the specified computer. Any preexisting NetBackup Administration Console and JRE packages are removed.

Hosts tab in Deployment management

Use the **Hosts** tab to indicate the hosts that you want to associate with the deployment policy.

To add hosts to a deployment policy

- 1
- Go to the **Hosts** tab.
- 2
- Select **Add hosts** or **Add**.
- The list of hosts that displays after you select **Add** or **Add hosts** are those hosts that are compatible with the package that you selected.
- If you see a warning icon besides a host name, it could be due to one of the following reasons:
- The selected package is missing for a particular operating system.
- The selected hosts are either at a lower or higher version than the selected package version. For Emergency binaries (EEBs), the versions must match.
- The host is already on the same version as the selected package.
- Information is not available for the host.
- 3
- From the list of hosts, select the hosts that you want to add to the deployment policy.
- 4
- Select **Add**.

Schedules tab in Deployment management

Use the **Schedules** tab in Deployment management for the following tasks:

- To view a summary of all schedules within that policy.
- To create a new schedule.
- To edit or delete an existing schedule.

The schedules that are defined on the **Schedules** tab determine when VxUpdate occur for the selected deployment policy. The calendar displays a summary of all the schedules.

The **Schedules** tab contains both schedule information and other configuration options, beyond when the job is to run.

Setting	Description
Name	Enter a name for the new schedule.

Setting	Description
Operation	Specify the type of operation that you want to associate with the schedule. Precheck - Performs the various precheck operations, including confirming there is sufficient space on the client for the update. The precheck schedule type does not exist for EEB packages. Stage - Moves the update package to the client, but does not install it. This operation also performs the precheck operation. Install - Installs the specified package. This operation also performs the precheck and the stage package operations. If you already performed the stage package operation, the install schedule does not move the package again. Note: Be aware that adding multiple different schedule types to the same deployment schedule window has unpredictable results. VxUpdate has no defined behavior to determine which schedule type runs first. If a single deployment schedule window has precheck, stage, and install jobs, there is no way to specify the order in which they run. The precheck or the stage schedules can fail, but the install completes successfully. If you plan to use precheck, stage, and install schedules, it is recommended that you create separate schedules and separate windows for each.
Start date	Specify the date and time you want the policy to start in the text field or with the date and the time spinner. You can also click the calendar icon and specify a date and time in the resulting window. You can select a schedule by clicking and dragging over the three-month calendar that is provided at the bottom of the window.
End date	Specify the date and time you want the policy to end as you specified the start time.

Security options tab in Deployment management

Use the policy **Security options** tab to configure the settings for external security certificates. These settings are only available if a selected host is configured to use external certificates (certificates that are signed by a CA other than the NetBackup CA).

Attribute	Description
Use existing certificates when possible	This option instructs NetBackup to use the existing NetBackup CA or external CA certificates, if available. By default, the Use existing certificates when possible option is selected. Deselecting the Use existing certificates when possible option lets you specify the location for external certificate authority information for both UNIX and Linux computers and Windows computers. Note: If you specify this option and certificates are not available, your upgrade fails.
From Windows certificate store (Only for Windows)	Specifies that the certificate from the Windows certificate store is used. The certificate is searched using the following details that are provided with the Certificate location: Store name, Issuer name, Subject name.
Certificate file	Specifies the path to the external certificate of the host.
Trust store location	Specifies the path to the <code>pem</code> bundle of the Certificate Authorities.
Private key file	Specifies the path to the private key for the external certificate of the host.
Passphrase file	Specifies the path to the text file where the passphrase for the external certificate's private key is stored.
CRL check level	Specifies the revocation check level for the external certificate. It also lets you disable the revocation check for the external certificates. Based on the check level, the status of the certificate is validated against the Certificate Revocation List (CRL) during host communication. You can choose to use the CRLs from the directory that is specified in the NetBackup configuration file or the CRL Distribution Point (CDP).
From certificate file path (for file-based certificates) (Only for Windows)	Specifies a list of comma-separated clauses where each clause element contains a query. The clause is of the form <code><Store name>\<Issuer Name>\<Subject Name></code>. <code>\$hostname</code> is a keyword that is replaced with the fully qualified domain name of the host. For certificate selection from the Windows certificate store, NetBackup can pick a certificate from any of the Local Machine certificate stores on a Windows host. ■ Store name – The certificate store where the certificate is present ■ Issuer name (optional) – The certificate issuer name ■ Subject name – The subject name of the certificate If the issuer name is not specified, the certificate is searched based on the subject name.

Copy a deployment policy

Use the **Copy policy** option to save time creating policies. This option is especially useful for the policies that contain many of the same policy attributes, schedules, or hosts selections.

To copy a deployment policy

- 1 Open the NetBackup web UI.
- 2 On the left, select **Hosts > Deployment management**.
- 3 Select the policy to copy.
- 4 Select **Copy policy**.
- 5 Enter the name for the new policy.
- 6 Select **Copy**. The only difference between the new policy and the copied policy is the name.

Make any changes that you want to the new policy. Then select **Copy**.

Manually deploy a deployment policy

You can manually initiate a deployment policy based on an existing policy. Manually initiate deployment policies when you are logged into the server locally and need to force an immediate update. Or you can initiate an immediate upgrade for emergency binaries.

Use the **Deploy now** option to initiate a deployment job manually.

To manually deploy a deployment policy

- 1 Open the NetBackup web UI.
- 2 Go to **Hosts > Deployment management**. Then click the **Deployment policies** tab.
- 3 Select the policy that you want to start, and select **Deploy now**.
- 4 Select the operation that you want to run and the hosts that you want to upgrade.

If you do not select any hosts, NetBackup upgrades all hosts.

- 5 Click **Deploy now** to start the manual deployment job.

Deployment job status

Monitor and review deployment job status in the Activity monitor. The **Deployment** job type is the new type for VxUpdate policies. Deployment policy parent jobs that exit with a status code 0 (zero) indicate that all the child jobs successfully completed. Parent jobs that finish with a status code 1 indicate that one or more of the child jobs succeeded, but at least one failed. Any other status code indicates failure. Review the status of the child jobs to determine why they failed. Otherwise, there are no differences between deployment jobs and other NetBackup jobs.

Your deployment job may receive a status code 224. This error indicates that the client's hardware and operating system are specified incorrectly. You can correct this error by modifying the deployment policy with the `bpplclients` command found in:

Linux: `/usr/opensv/netbackup/bin/admincmd`

Window: `install_path\netbackup\bin\admincmd`.

Use the syntax shown:

```
bpplclients deployment_policy_name -modify client_to_update -hardware  
new_hardware_value -os new_os_value
```

Deployment policies use a simplified naming scheme for operating system and hardware values. Use the values as shown for the `bpplclients` command:

Table 9-1 Deployment policy operating system and hardware

Operating system	Hardware
debian	x64
redhat	x64
suse	x64
redhat	ppc64le
suse	ppc64le
redhat	zseries
suse	zseries
aix	rs6000
solaris	sparc
solaris	x64

Table 9-1 Deployment policy operating system and hardware *(continued)*

Operating system	Hardware
windows	x64

Security certificates are not deployed as part of the VxUpdate upgrade if the **Security level for certificate deployment** is set to **Very High**. This setting is located in the **Global security** settings.

See “[Select a security level for NetBackup certificate deployment](#)” on page 478.

If you cannot communicate with your clients after you use VxUpdate to upgrade your clients, please ensure that the proper security certificates were issued during upgrade. You may need to manually deploy the certificates. Refer to the following article that is shown for additional details:

https://www.veritas.com/content/support/en_US/article.100039650

Configuring storage

- [Chapter 10. Overview of storage options](#)
- [Chapter 11. Configuring disk storage](#)
- [Chapter 12. Managing media servers](#)
- [Chapter 13. Configuring storage units](#)
- [Chapter 14. Configuring robots and tape drives](#)
- [Chapter 15. Configuring tape media](#)
- [Chapter 16. Inventorying robots](#)
- [Chapter 17. Staging backups](#)
- [Chapter 18. Troubleshooting storage configuration](#)

Overview of storage options

This chapter includes the following topics:

- [About storage configuration](#)

About storage configuration

NetBackup lets you configure storage options for all protection plans and policies. To set up storage options, on the left click **Storage**.

You can configure the following types of storage:

- Disk storage
 - See “[Disk storage unit considerations](#)” on page 273.
 - See “[Integrating MSDP Cloud and CMS](#)” on page 245.
 - See “[Create an AdvancedDisk, OpenStorage \(OST\), or Cloud Connector storage server](#)” on page 250.
 - For details on Universal Shares, see the [NetBackup Deduplication Guide](#).
 - For details on Cloud connector, see the [NetBackup for Cloud Object Store Administrator's Guide](#).
 - See “[About the MSDP object store](#)” on page 258.
- Media servers
 - See “[Add a media server](#)” on page 261.
- Snapshot Manager
 - See the [NetBackup Snapshot Manager for Data Center Administrator's Guide](#) for details.
- Storage lifecycle policies (SLPs)
 - See the [NetBackup Administrator's Guide, Volume I](#).

- Storage units
See [“Overview of storage units”](#) on page 266.
- Tape storage
See [“Configure drives and robots by using the wizard”](#) on page 281.

Note: If you use Key Management Service (KMS), it must be configured before you can select the KMS option in the storage server setup. Refer to [NetBackup Security and Encryption Guide](#) for more information.

To ensure that A.I.R. and other storage capabilities are displayed accurately for the storage servers on the NetBackup web UI, upgrade the media server. You must upgrade the media server that has NetBackup versions 8.2 or earlier. After you upgrade the media server then use the command line to update the storage server.

Use the following command to update the storage server:

```
/usr/opensv/netbackup/bin/admincmd/nbdevconfig -updatests  
-storage_server <storage server name> -stype PureDisk
```

For more information, refer to the [NetBackup Deduplication Guide](#).

Configuring disk storage

This chapter includes the following topics:

- [Create a Media Server Deduplication Pool storage server](#)
- [Integrating MSDP Cloud and CMS](#)
- [Create a Media Server Deduplication Pool \(MSDP\) storage server for image sharing](#)
- [Create an AdvancedDisk, OpenStorage \(OST\), or Cloud Connector storage server](#)
- [Create an MSDP server for MSDP volume group \(MVG\)](#)
- [Create the MVG volume](#)
- [Edit a storage server](#)
- [About configuring disk pool storage](#)
- [Share images from an on-premises location to the cloud](#)
- [Overview of universal shares](#)
- [About the MSDP object store](#)

Create a Media Server Deduplication Pool storage server

Use this procedure to create a Media Server Deduplication Pool storage server. You have the option to create a disk pool (local storage or cloud storage) and storage unit after you create a storage server. The recommendation is that you create the disk pool and storage unit if they do not exist in NetBackup.

To add an MSDP storage server

- 1 On the left, select **Storage > Disk storage**. Select the **Storage servers** tab, then click **Add**.
- 2 In the **Storage type** list, select **Disk storage servers**.
- 3 From the **Category** options, select **Media Server Deduplication Pool (MSDP, MSDP Cloud, MVG)**.
Click **Start**.
- 4 In **Basic properties**, enter all the required information.

To select your media server, click the search icon. If you do not see the media server you want to use, you can use **Search** field to find it.

Click **Next**.
- 5 In the **Storage server options**, enter all required information and click **Next**.

If you use Key Management Service (KMS), it must be configured before you can select the **KMS** option.
- 6 (Optional) In **Media servers**, click **Add** to add any additional media servers that you want to use.

Click **Next**.
- 7 On the **Review** page, confirm that all options are correct and click **Save**.

If the MSDP storage server creation is unsuccessful, follow the prompts on the screen to correct the issue.

To configure MSDP to use cloud storage, use the following procedure (drop-down in **Volumes** step) to select an existing disk pool volume or create a new one.
- 8 (Optional) At the top, click on **Create disk pool**.
- 9 (Optional) To create a cloud logical storage unit and disk pool with replication, click on **Create disk pool**.

Enter the required information to create a disk pool.

In the next tab, select and add the required cloud volume. Select the cloud storage provider and the required details of the storage provider. Enter the credentials to access the cloud storage provider and then define the advanced settings.

For the cloud logical storage unit, click **Edit** to update the **Cloud cache properties** setting in the corresponding disk pool properties page. You must restart the `pdde` services for the updated setting to work.

Additional notes

Review the following additional information:

- Currently, AWS S3 and Azure storage API types are supported.
For more information about the storage API types that NetBackup supports, refer to the topic *About the cloud storage vendors for NetBackup* in the [NetBackup Cloud Administrator's Guide](#).
- When you enable Server-Side Encryption, you can configure AWS Customer-Managed keys. These keys cannot be deleted once they are in use by NetBackup. Each object is encrypted with the key during upload and deleting the key from AWS causes NetBackup restore failures.
- For more information on environments and deployment of Veritas Alta Recovery Vault for NetBackup, refer to the following article:
<https://www.veritas.com/docs/100051821>
Before you enable the Veritas Alta Recovery Vault Azure and Azure Government options, review the steps from the *Configuring Veritas Alta Recovery Vault Azure and Azure Government* section in the [NetBackup Deduplication Guide](#).
Veritas Alta Recovery Vault supports multiple options. For Veritas Alta Recovery Vault Azure and Azure Government options in the web UI, you must contact your Cohesity NetBackup account manager for credentials or with any questions.

Integrating MSDP Cloud and CMS

Note: CMS is now supported for all S3 and Azure cloud vendor types.

To integrate MSDP Cloud and CMS

- 1 If you haven't already, create an MSDP storage server. See the *Configuring MSDP server-side deduplication* section in the *NetBackup Deduplication Guide*.
- 2 Add a disk pool.

On the left, select **Storage > Disk storage** and select the **Disk pools** tab. Then select **Add**.
- 3 In the **Disk pool options**, click **Change** to select a storage server.
 - Select a storage server from the list and click **Select**.
 - Enter the **Disk pool name**.
 - If **Limit I/O streams** is cleared, the default value is Unlimited and may cause performance issues.
 - After all the required information is added, click **Next**.

- 4 In the **Volumes** properties, from the **Volume** list select **Add volume**.
 - Provide a unique volume name that gives adequate description of the volume.
 - For the **Cloud storage provider**, select Microsoft Azure, Amazon, or any other cloud provider of S3 and Azure types. Then click **Select**.
- 5 In the **Region** section, select the appropriate region.
- 6 In the **Associate Credentials** section, select an authentication type, select **Add a New Credential**.

Enter a **Credential name** which should be a valid name and should only contain alphanumeric characters, hyphen, colon, and underscore.

Note: For details of authentication types like AWS IAM Role Anywhere and Azure Service Principal, see the *NetBackup Deduplication Guide*.

- 7 In **Access details for type account**, select **AWS S3 compatible** or **Azure Blob** and enter the access information.

Alternatively, you can use **Select existing credential** but the credentials must have a Category of MSDP-C and proper credentials for the chosen supported cloud provider.
- 8 In the **Cloud buckets** section, select from the following options:
 - If the cloud credentials in use do not have the permissions to list buckets, click **Enter an existing cloud bucket name**.
 - To create a cloud bucket, click **Select or add a cloud bucket**. Then click **Retrieve list** to select a predefined bucket from the list.
- 9 Click **Next**.
- 10 In **Replication**, click **Next**.
- 11 On the **Details** page, verify that all settings and information are correct. Click **Finish**.

The disk pool creation and replication configuration continue in the background if you close the window. If there is an issue with validating the credentials and configuration of the replication, you can use the **Change** option to adjust any settings.

In the **Volumes** step, you can now use **Retrieve List** (list buckets) or create a bucket depending on what you want to accomplish.

Update the credentials

To update the credentials

- 1 Create a disk pool.
- 2 After you have selected **Add volume**, **Volume name**, select **Cloud storage**, and select a **Region** then click **Select existing credential**.
- 3 Locate **Credential name**. Then click **Actions > Edit**.
- 4 Make any changes as necessary.
- 5 In the **Permissions**, make any changes as necessary and click **Save**.
- 6 Finish adding the disk pool.

nbcdutil changes

- (10.3 and later) Use the parameter `cmscredname` instead of `username`. However, `username` is still supported for older media servers.
- Validate credentials. `nbcdutil -validatecreds -storage_server mystorage_server -cmscredname mycmscredentialname`
- Create a bucket. `nbcdutil -createbucket -storage_server mystorage_server -cmscredname mycmscredentialname -bucket_name bucketname`

nbdevconfig changes

- You need to provide `lsuCmsCredName` in the configuration file for Veritas Alta Recovery Vault Azure and Veritas Alta Recovery Vault Azure Gov.
- Instead of using the storage account name for `lsuCmsCredName`, use the name of the credentials that are created when you use **Credential management**.
- The configuration file for `nbdevconfig` CLI now uses a new Key `cmsCredName` instead of user `lsuCloudUser` and `lsuCloudPassword`. The file should look like the following:

```
[root@vramsingh7134 openv]# cat /add_lsu.txt
V7.5 "operation" "add-lsu-cloud" string
V7.5 "lsuName" "ms-lsu-cli" string
V7.5 "lsuCloudBucketName" "ms-mybucket-cli" string
V7.5 "lsuCloudBucketSubName" "ms-lsu-cli" string
V7.5 "cmsCredName" "aws-creds" string
V7.5 "requestCloudCacheCapacity" "4" string
```

Note: For regular Azure and AWS from this 10.3 and newer: If you use the `createdv` option to create a cloud bucket on the primary server or media server or on an older media server, you see a message that tells you to use `nbcldutil`.

Note: Some browsers like Firefox may auto-populate the fields to store the credentials in the CMS with credentials the browser saves. You must to turn off a setting in Firefox so that the credentials do not auto-populate.

Migrating or updating MSDP Cloud and CMS

You can update only Access key credentials to CMS. You cannot update the configured credentials for an older disk pool to CMS to use other authentication types. Upgraded credentials to use in CMS must be Access key based.

To migrate or update MSDP Cloud

- 1 If using an MSDP on an old NetBackup version, configure MSDP cloud for any cloud provider by providing credentials in the **Access details for account** section.
- 2 Run a backup and restore.
- 3 Upgrade the MSDP to the newest version.
- 4 Click on the MSDP cloud disk pool that was configured in the previous release.
- 5 In the **Associate credentials** box, select **Actions > Replace**. Or, to update the credentials for the disk pool, select **Edit**.
- 6 Select **Continue**.
- 7 Provide the appropriate credentials and select **Next**.
- 8 Follow the steps from Credential Management.
- 9 Select **Save**.
- 10 Restart the NetBackup services on the primary server and the media server.

Create a Media Server Deduplication Pool (MSDP) storage server for image sharing

Use this topic to create a cloud recovery server for image sharing. Refer to the *About image sharing using MSDP cloud* topic in the [NetBackup Deduplication Guide](#) for more information about a cloud recovery server.

To configure cloud recovery server:

- 1** On the left, click **Storage > Disk storage**. Click the **Storage servers** tab, then click **Add**.
- 2** In the Storage type drop-down, select the option you want to use.
- 3** Select **Media Server Deduplication Pool (MSDP) for image sharing** from the list.
- 4** In the **Basic properties**, enter all the required information and click **Next**.
You must select your media server by clicking on the field. If you do not see the media server you want to use, use the search option.
- 5** In the storage server options, enter all the required information except for **Encryption options** and **Encryption for local storage** and click **Next**.
If KMS encryption is enabled for the on-premises side, Key Management Service (KMS) must be configured before you can configure cloud recovery server. In the cloud recovery host, you must not configure KMS encryption when you set up a storage server. The KMS options from the on-premises side are selected and configured automatically in the cloud recovery host.
- 6** (Optional) In Media servers, click **Next**. As the cloud recovery server is an all-in-one NetBackup server, no additional media servers are added.
- 7** On the **Review** page, confirm that all options are correct and click **Save**.
If the MSDP with the image sharing creation is unsuccessful, follow the prompts on the screen to correct the issue.
- 8** At the top, click on **Create disk pool**.
You can also create a disk pool as follows:
On the left, click **Disk storage**. Click the **Disk pools** tab, then click **Add**.
- 9** In **Disk pool** options, enter all the required information and click **Next**.
Click **Change** to select a storage server.
- 10** In **Volumes**, use the **Volume** drop down to add a new volume. Enter all the required information based on the selection and click **Next**.
The volume name must be same as the volume name that is on the on-premises side or the sub bucket name.
The storage class must be same as the storage class that is selected on the on-premises side. If you do not select the correct class, imports or restores may fail later.

- 11 In **Replication**, click **Next** to continue without adding any primary server.
- 12 On the **Review** page, verify that all settings and information are correct. Click **Save**.

Note: In an image sharing all-in-one setup, the **Add** button is not available on the disk pool page if the image sharing server already has an associated disk pool.

See [“Share images from an on-premises location to the cloud”](#) on page 257.

Create an AdvancedDisk, OpenStorage (OST), or Cloud Connector storage server

Use the following procedures to create AdvancedDisk, OpenStorage, or a Cloud Connector storage server.

Create an AdvancedDisk storage server

Follow this procedure to create an AdvancedDisk storage server.

To create an AdvancedDisk storage server

- 1 On the left, select **Storage > Disk storage**. Select the **Storage servers** tab, then select **Add**.
- 2 In the **Storage type** list, select **Disk storage servers**.
- 3 From the **Category** options, select **AdvancedDisk**.
- 4 Select **Start**.
- 5 Select a media server from the list and select **Select**.

Create an OpenStorage (OST) storage server

Follow this procedure to create an OpenStorage (OST) storage server.

To create an OpenStorage (OST) storage server

- 1 On the left, select **Storage > Disk storage**. Select the **Storage servers** tab, then select **Add**.
- 2 In the **Storage type** list, select **Disk storage servers**.
- 3 From the **Category** options, select **OpenStorage (OST)**.
- 4 Select **Start**.

- 5** In **Basic properties**, enter all the required information.
- To select your media server, select the search icon. If you do not see the media server you want to use, you can use **Search** field to find it.
- Select the correct **Storage server type**.
- Select **Next**.
- 6** (Optional) In **Media servers**, select **Add** to add any additional media servers you want to use.
- Select **Next**.
- 7** On the **Review** page, confirm that all options are correct and select **Save**.
- After you select **Save**, the credentials you entered are validated. If the credentials are invalid, select **Change** and you can correct the issue with the credentials.
- 8** (Optional) At the top, select **Create disk pool**.

Create a Cloud Connector server

Follow this procedure to create a Cloud storage server.

To create a Cloud storage server

- 1** On the left, select **Storage > Disk storage**. Select the **Storage servers** tab, then select **Add**.
- 2** In the **Storage type** list, select **Disk storage servers**.
- 3** In the **Storage type** list, select **Cloud connector**.
- 4** Select **Start**.
- 5** In **Basic properties**, enter all the required information.
- You must select your **Cloud storage provider** by selecting on the field. If you do not see the cloud storage provider you want to use, you can use **Search** to find it.
- If the **Region** information that you want to select does not appear in the table, use **Add** to manually add the required information. This option does not appear for every cloud storage provider.
- To select your media server, select the search icon. If you do not see the media server you want to use, you can use **Search** field to find it.
- Select **Next**.

- 6 In **Access settings** enter the required access details for the selected cloud provider and select **Next**.

If you use `SOCKS4`, `SOCKS5`, or `SOCKS4A`, some of the options in the **Advanced** section are not available.
- 7 In **Storage server options**, you can adjust the **Object size**, enable compression, or encrypt data and then select **Next**.
- 8 (Optional) In **Media servers**, select **Add** to add any additional media servers you want to use.

For Cloud storage servers, media servers with a NetBackup version that is older than the primary server are not listed.

Select **Next**.
- 9 On the **Review** page, confirm that all options are correct and select **Save**.
- 10 (Optional) At the top, select **Create disk pool**.

Create an MSDP server for MSDP volume group (MVG)

The MSDP server that has the MVG capability and manages the MVG volumes is called as MVG server. For more information about MSDP volume group (MVG) see *NetBackup Deduplication Guide*.

To create an MSDP server for MSDP volume group

- 1 On the left, click **Storage > Disk storage**. Click the **Storage servers** tab, then click **Add**.
- 2 In the **Storage type** list, select **Disk storage servers**.
- 3 From the **Category** options, select **Media Server Deduplication Pool (MSDP, MSDP Cloud, MVG)**.

Click **Start**.
- 4 In **Basic properties**, enter all the required information.

To select your media server, click the search icon. If you do not see the media server you want to use, you can use **Search** field to find it.

Click **Next**.

- 5 In the **Storage server** options, enter all required information and select **Enable MSDP volume group (MVG) service**.

This option configures an MSDP server as an MSDP volume group (MVG) server, which lets you group the volumes from other MSDP servers to create MVG volumes. After you enable it, the MVG server can only host MVG volumes and cannot host its own local or cloud volumes.

If you use Key Management Service (KMS), it must be configured before you can select the **KMS** option.

Click **Next**.

- 6 (Optional) In **Media servers**, click **Add** to add any additional media servers that you want to use.

Click **Next**.

- 7 On the **Review** page, confirm that all options are correct and click **Save**.

Create the MVG volume

The MSDP volume group (MVG) is an MSDP capability to build the volume groups in the storage layer on top of individual MSDP storage servers. One volume group is represented to NetBackup as a virtual volume, and this virtual volume is called MVG volume.

For more information about MSDP volume group (MVG) see *NetBackup Deduplication Guide*.

To create the MVG volume

- 1 On the left, click **Storage > Disk storage**.
- 2 Click **Storage > Disk storage**. Click the **Disk pools** tab, then click **Add**.
- 3 In the **Disk pool options**, click **Change** to select a storage server that is MVG enabled.
- 4 Enter all the required information. Click **Next**.
- 5 In **Volumes**, from the **Volume** drop down, select **Add MVG volume**.
- 6 Enter the MVG volume name and select the attributes to filter the desired volumes.
- 7 Select multiple volumes for the MVG volume. Click **Next**.

- 8 On the **Review** page, verify that all settings and information are correct. Click **Finish**.
- 9 After MVG volume is added, the MVG column appears under the **Disk pools** tab.

Use the disk pool of the MVG volume in the same way as before to configure storage unit and replication targets.

Edit a storage server

This procedure tells you how to edit a storage server.

To edit a storage server

- 1 On the left, click **Storage > Disk storage**.
- 2 Click the **Storage servers** tab.
- 3 Click the name of the storage server that you want to edit.
- 4 On the storage server review page, locate **Troubleshooting properties**. Then select **Edit**.
- 5 Under **Universal share properties** click **Edit** to edit the universal share properties.
- 6 Under **Media servers** click **Add** to add the load-balancing media servers.
For more information see the *Adding an MSDP load-balancing server* topic in the *NetBackup Deduplication Guide*.
- 7 Under *Isolated recovery environment*, you can configure isolated recovery environment on the storage server if required.

For more information see the topic *Configuring an isolated recovery environment using the web UI* in the *NetBackup Deduplication Guide*.

About configuring disk pool storage

You can configure disk pools if you license a NetBackup feature that uses disk pools.

For more information, see the following guides:

- The *NetBackup AdvancedDisk Storage Solutions Guide*.
- The *NetBackup Cloud Administrator's Guide*.
- The *NetBackup Deduplication Guide*.
- The *NetBackup OpenStorage Solutions Guide for Disk*.

- The *NetBackup Replication Director Solutions Guide*.

Create a disk pool

Use this procedure to create a disk pool after you create any type of storage server. You can create a disk pool at any time, but to create a disk pool requires that you have an existing storage server created.

When you view the **Disk pools** tab, the **Available space** column can be empty for a disk pool that uses a cloud storage provider. NetBackup cannot retrieve the information because the cloud provider does not supply an API for that information.

To create a disk pool

- 1 On the left, click **Storage > Disk storage**. Click the **Disk pools** tab, then click **Add**.

Another way to create a disk pool is to click **Create disk pool** at the top of the screen after you have created a storage server.

- 2 In the **Disk pool options**, enter all the required information.

Click **Change** to select a storage server.

If the option **Limit I/O streams** is cleared, the default value is **Unlimited** and may cause performance issues.

Click **Next**.

- 3** In **Volumes**, use the **Volume** drop down to select a volume or add a new volume. If you want to add a new disk pool volume, use the **Add volume** option.

You can configure an MSDP storage server to use cloud storage. Select an existing cloud volume or create a new volume for the MSDP storage server.

Note: When you enable Server-Side Encryption, you can configure AWS Customer-Managed keys. These keys cannot be deleted once they are in use by NetBackup. Each object is encrypted with the key during upload and deleting the key from AWS causes NetBackup restore failures.

Note: Veritas Alta Recovery Vault supports multiple options. For Veritas Alta Recovery Vault Amazon and Amazon Government options in the web UI, you must contact your Cohesity NetBackup account manager for credentials or with any questions.

For more information on environments and deployment, refer to [Cohesity Alta Recovery Vault](#).

For more information about Cohesity Alta Recovery Vault Azure options, refer to *About Cohesity Alta Recovery Vault Azure* in the [NetBackup Deduplication Guide](#).

Enter all the required information based on the selection and click **Next**.

- 4** In **Replication**, click **Add** to add replication targets to the disk pool.

This step lets you select a trusted primary server or add a trusted primary server. You can add a primary server that supports NetBackup Certificate Authority (NBCA), ECA, and ECA together with NBCA.

Replication is supported only on MSDP.

Review all the information that is entered for the replication targets and then click **Next**.

- 5** On the **Review** page, verify that all settings and information are correct. Click **Finish**.

The disk pool creation and replication configuration continue in the background if you close the window. If there is an issue with validating the credentials and configuration of the replication, you can use the **Change** option to adjust any settings.

Note: The image sharing server, which already has a disk pool that is created for it, is not available while NetBackup creates a new disk pool for image sharing.

Edit a disk pool

This procedure describes how to edit a disk pool.

To edit a disk pool

- 1 On the left, select **Storage > Disk storage**. Select the **Disk pools** tab.
- 2 Click on the name of the disk pool that you want to edit.
- 3 On the disk pool details page, select **Edit** to edit the parameters of the disk pool.
- 4 (MSDP) Under **Replication targets**, select the **Add** button to add the replication targets.

Share images from an on-premises location to the cloud

You can share images from an on-premises location to the cloud. Set up a cloud recovery server on demand and then share the images to that server.

Use the information from the following topic from the [NetBackup Deduplication Guide](#) to set up a cloud recovery server: *About image sharing using MSDP cloud*.

Steps to complete after setting up the cloud recovery server

Before you begin, ensure that you have the required permissions in the web UI to import the image, restore, convert, and access the AMI ID or VHD.

To import the images

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Disk storage**. Then click the **Disk pools** tab.
- 3 Select the volume pools that contain the images that you want to share.
- 4 In the Disk pool options, locate the disk pool name and click **Actions > Fast Import**.

Note: The fast import option is an import operation that is specific to image sharing. You can import the backed-up images from the cloud storage to the cloud recovery server that is used for image sharing. After a fast import, you can restore the images. For AWS cloud provider, you can also convert the VM image to an AWS AMI. For Azure cloud provider, you can convert the VM image to VHD.

- 5 In the **Fast import images** page, select the backup images that you want to import and click **Import**.
- 6 Verify the activity completion status in the **Activity monitor**.

To convert the VM images to AWS AMI or VHD in Azure

- 1 On the left, click **Workloads > VMware**. Then select the imported VMware image to convert.
- 2 On the **Recovery point** tab, select the recovery date.
- 3 For the recovery point date, choose the required recovery point. Click **Actions > Convert**.

For Veritas Alta Recovery Vault, it may take time to get the disk volume and the credentials information.

Provide the credentials of Azure general-purpose storage accounts or AWS account with IAM and EC2 related permissions.

For more information on the permission, see *Recover the VM as an AWS EC2 AMI or VHD in Azure* topic of the *NetBackup Deduplication Guide*.

- 4 After the conversion is complete, an AMI ID or VHD URL is generated.
- 5 Use the AMI ID to locate the image in AWS and then use the AWS console to start the EC2 instance. Or use VHD URL to create a virtual machine.

Overview of universal shares

The universal share feature provides data ingest into an existing NetBackup deduplication pool (MSDP) or a supported Cohesity appliance using an NFS or a CIFS (SMB) share.

Both universal shares and MSDP use deduplication and compression.

Space efficiency is achieved by storing this data directly into an existing NetBackup-based Media Server Deduplication Pool.

For more information about universal shares, refer to the following guide:

[NetBackup Deduplication Guide](#)

About the MSDP object store

S3 interface for MSDP provides S3 APIs in MSDP server. You must configure the MSDP object store to use S3 interface for MSDP.

S3 interface for MSDP is compatible with Amazon S3 cloud storage service. It supports most of the commonly used S3 APIs such as create bucket, delete bucket, store object, retrieve object, list object, delete object, multipart upload, and so on.

S3 interface for MSDP also supports object versioning, IAM, and identity-based policy. It uses snowball-auto-extract to support small objects batch upload.

See [“Configuring the MSDP object store”](#) on page 259.

See [“Resetting the MSDP object store root user credentials”](#) on page 259.

Configuring the MSDP object store

You can configure MSDP object store on the MSDP build-your-own (BYO) and Flex appliance platform using the NetBackup web UI. See [About S3 Interface for MSDP](#) topic in the *NetBackup Deduplication Guide* to configure MSDP object store on other platforms.

To configure the MSDP object store

- 1 Configure an MSDP storage server if required.
See [“Create a Media Server Deduplication Pool storage server”](#) on page 243.
- 2 On the left, click **Storage > Disk storage**.
- 3 On the **MSDP object store** tab, click **Add**.
- 4 Provide the following required information:
 - Select the storage server.
 - Select the certificate authority type. Only NBCA (NetBackup certificate authority) certificate is supported in the web UI.
 - Enter the port number for the MSDP S3 interface. The default port number is 8443.
- 5 Click **Save** and wait for the response.
S3 interface for MSDP starts and S3 server root user credentials are generated. A screen appears with the root user access key, secret key, and MSDP S3 service endpoint. You must store the keys and endpoint manually.

Resetting the MSDP object store root user credentials

You can reset the root user credentials for the MSDP object store endpoint that is added on the NetBackup web UI.

To reset MSDP object store root user credentials

- 1** On the left, click **Storage > Disk storage**. Click the **MSDP object store** tab.
- 2** Locate the MSDP object store endpoint and click **Reset IAM root** at the right.
- 3** S3 server root user credentials are reset.

A screen appears with root user access key, secret key, and MSDP S3 service endpoint. You must store the keys and endpoint manually.

Managing media servers

This chapter includes the following topics:

- [Add a media server](#)
- [Activate or deactivate a media server](#)
- [Stop or restart the media device manager](#)
- [About NetBackup server groups](#)
- [Add a server group](#)
- [Delete a server group](#)

Add a media server

The following table describes an overview of how to add a media server to an existing NetBackup environment.

Note: The NetBackup Enterprise Media Manager service must be active when a media server is added, devices and volumes are configured, and clients are backed up or restored.

Table 12-1 Adding a media server

Step	Procedure	Section
Step 1	On the new media server host, attach the devices and install any software that is required to drive the storage devices.	See the vendor's documentation.
Step 2	On the new media server host, prepare the host's operating system.	See the NetBackup Device Configuration Guide .

Table 12-1 Adding a media server (*continued*)

Step	Procedure	Section
Step 3	On the primary server, add the new media server to the Media servers list of the primary server. Also, add the new media server to the Additional servers list of the clients that the new media server backs up. If the new media server is part of a server group, add it to the Additional servers list on all media servers in the group. Note: Ensure that the name you use in NetBackup is the same as the host name in the TCP/IP configuration.	See the <i>Servers properties</i> topic in the NetBackup Administrator's Guide, Volume I .
Step 4	Install the NetBackup media server software on the new host.	See the NetBackup Installation Guide .
Step 5	On the primary server, configure the robots and drives that are attached to the media server.	See the <i>Configuring robots and tape drives by using the wizard</i> topic in the NetBackup Administrator's Guide, Volume I .
Step 6	On the primary server, configure the volumes.	See the <i>About adding volumes</i> topic in the NetBackup Administrator's Guide, Volume I .
Step 7	On the primary server, add storage units to the media server. Always specify the media server as the media server for the storage unit.	See "Create a storage unit" on page 268.
Step 8	On the primary server, configure the NetBackup policies and schedules to use the storage units that are configured on the media server.	See "Add a policy" on page 357.
Step 9	Test the configuration by performing a user backup or a manual backup that uses a schedule that specifies a storage unit on the media server.	See "Perform manual backups" on page 365.

Activate or deactivate a media server

When you activate a media server, NetBackup can use it for backup and restore jobs. You can deactivate a media server. A common reason to do so is to perform maintenance. When a media server is deactivated, NetBackup does not send job requests to it.

When you deactivate a media server, the following things occur:

- Current jobs are allowed to complete.
- If the host is part of a shared drive configuration, it does not scan drives.

To activate or deactivate a media server

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Media servers**. Then click the **Media servers** tab.
- 3 Select the media server to activate or deactivate.
- 4 Click **Activate** or **Deactivate**.

Stop or restart the media device manager

Use the following procedure to stop and restart the NetBackup device manager.

To start or stop the media device manager

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Media servers**. Then select the **Media servers** tab.
- 3 Select the media server and select **Stop/Restart media manager device daemon**.
- 4 Locate **Action** and select the action that you want to take.

The actions that are available depend on the state of the media manager device.
- 5 Select any of the **Options** that you want.
- 6 Click **Apply**.

About NetBackup server groups

A server group is a group of NetBackup servers that are used for a common purpose.

A NetBackup **Media sharing** group is a server group that shares tape media for write purposes (backups). All members of a **Media sharing** server group must have the same NetBackup primary server.

A **Media sharing** group can contain the following:

- NetBackup primary server
- NetBackup media servers
- NDMP tape servers

Add a server group

A server group is a group of NetBackup servers that are used for a common purpose. Servers can be in more than one group.

Caution: NetBackup allows a server group name to be the same as the name of a media server. To avoid confusion, do not use same name for a server group and a media server.

To add a server group

- 1 On the left, click **Storage > Media servers**.
- 2 Click **Server groups**.
- 3 Click **Add server group**.
- 4 Provide the information for the server group.

Server group name	Provide a unique name for the server group. Do not use the name for an existing media server or other host. You cannot change the name of an existing server group.
Server group type	Select the type of server group.
State	Active. The server group is available for use. Inactive. The server group is not available for use.
Description	Provide a description of the group.

- 5 To add a server to the group, click **Add**, select the server, then click **Add**.
To remove a server from the group, select the server and click **Remove**.
- 6 Click **Save**.

Delete a server group

You can delete a server group if it is no longer in use. Or, if the purpose of the servers in the group has changed.

To delete a server group

- 1** On the left, click **Storage > Media servers**.
- 2** Click **Server groups**.
- 3** Select the group to delete. Then click **Delete > Delete**.

Configuring storage units

This chapter includes the following topics:

- [Overview of storage units](#)
- [About configuring BasicDisk storage](#)
- [Create a storage unit](#)
- [Edit storage unit settings](#)
- [Copy a storage unit](#)
- [Delete a storage unit](#)
- [Tape storage unit considerations](#)
- [Disk storage unit considerations](#)
- [NDMP storage unit considerations](#)

Overview of storage units

The data that is generated from a NetBackup job is recorded into a type of storage that NetBackup recognizes. NetBackup recognizes the following storage configurations, all of which are configured in **Storage**.

Storage units

A storage unit is a label that NetBackup associates with physical storage or cloud storage. The label can identify a path to a volume or a disk pool. Storage units can be included as part of a storage lifecycle policy.

The following types of storage units are available in the NetBackup web UI.

Table 13-1 Storage unit types

Storage type	Storage unit type	Storage type or location	Option required
Drives and robots	Tape (Media manager)	Points to a robot or a standalone drive	
	NDMP	Points to an NDMP host (NDMP option)	NDMP option
Disk storage servers	Media Server Deduplication Pool (MSDP)	Points to the local or the cloud storage.	Data Protection Optimization Option
	AdvancedDisk	Points to a disk pool (storage directly attached to a media server).	Data Protection Optimization Option
	OpenStorage Technology (OST)	Points to a disk pool of the type <i>StorageName</i> .	OpenStorage Disk Option
	Cloud Connector	Points to a disk pool of the type <i>VendorName</i> , where <i>VendorName</i> can be the name of a cloud storage provider.	
	BasicDisk	Points to a directory.	

Storage unit groups

Storage unit groups let you identify multiple storage units as belonging to a single group. The NetBackup administrator configures how the storage units are selected within the group when a backup or a snapshot job runs.

Storage lifecycle policies

Storage lifecycle policies let the administrator create a storage plan for all of the data in a backup or snapshot.

About configuring BasicDisk storage

A **BasicDisk** type storage unit consists of a directory on locally-attached disk or network-attached disk. The disk storage is exposed as a file system to a NetBackup media server. NetBackup stores backup data in the specified directory.

No special configuration is required for **BasicDisk** storage. You specify the directory for the storage when you configure the storage unit.

Create a storage unit

Use this procedure to create a storage unit. You should create a storage unit after you create any type of storage server and disk pool. You can also follow these steps if you create a new storage unit without creating a storage server and disk pool.

When you view the **Storage units** tab, the **Used space** column may be empty for a storage unit that uses a cloud storage provider. NetBackup cannot retrieve the information because the cloud provider does not supply an API for that information.

See [the section called “Create a storage unit for a disk storage server”](#) on page 268.

See [the section called “Create a tape storage unit”](#) on page 269.

Create a storage unit for a disk storage server

To create a storage unit for a disk storage server

- 1 On the left, select **Storage > Storage units**. Select the **Storage units** tab, then select the **Add** button.

Another way to create a storage unit is to select the **Create storage unit** button at the top of the screen after you have created a disk pool.

- 2 For the **Storage type** list, select the option **Disk storage servers**.
- 3 Select the storage unit **Category** from the list and select **Start**.
- 4 In **Basic properties**, enter all the required information and click **Next**.
- 5 In **Disk pool**, select the disk pool you want to use in the storage unit and then select the **Next** button.

The **Enable WORM** option is activated when you select a disk pool that supports WORM (Write Once Read Many) storage.

For more information about WORM properties, refer to the topic *Configuring immutability and indelibility of data in NetBackup* in the [NetBackup Administrator's Guide, Volume I](#).

The **On demand only** option specifies whether the storage unit is available exclusively on demand. A policy or schedule must be explicitly configured to use this storage unit.

- 6 On the **Media server** tab, select the media servers you want to use and then select the **Next** button. Choose from the following options.

- **Allow NetBackup to automatically select**
NetBackup automatically selects the media server to use.
 - **Manually select**
Select the specific media servers that you want to use.
- 7 Select the **Next** button.
 - 8 Review the setup of the storage unit and then select the **Save** button.

Create a tape storage unit

To create a tape storage unit

- 1 On the left, click **Storage > Storage units**. Select **Add**.
- 2 In the **Storage type** list, select the option **Drives and robots** and select **Start**.
- 3 In the **Basic properties**, enter all the required information and select the **Next** button.
- 4 In the **Storage devices**, select the appropriate storage device and select the **Next** button.
- 5 In the **Media server**, the media servers are listed based on the storage device that you selected. Choose to allow NetBackup to automatically select the media server. Or, manually select the media server. Select the **Next** button.
- 6 In the **Review**, verify all the selections. You can also edit the details if any changes are required and select the **Save** button.

Note: The image sharing disk pool is not available as NetBackup creates a new storage unit.

More information

See [“Create a disk pool”](#) on page 255.

See [“Create a Media Server Deduplication Pool storage server”](#) on page 243.

See [“Create an AdvancedDisk, OpenStorage \(OST\), or Cloud Connector storage server”](#) on page 250.

See [“Create a protection plan”](#) on page 368.

Edit storage unit settings

This option is only available for the **Disk** storage unit type. You can edit the settings for the storage type **Disk storage servers**, but not **Drives and robots**.

Only make changes to a storage unit during periods when no backup activity is expected. This way backups are not affected for the policies or protection plans that use the affected storage units.

To edit storage unit settings

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Storage units**. Select the **Storage units** tab.
- 3 Select the storage unit that you want to edit.
- 4 Click **Edit** and make the required changes.

For example, you can edit the following settings:

- The basic properties of the storage unit.
- Additional properties
- Media servers
- Staging schedule

To edit a tape storage unit

- 1 On the left, click **Storage > Storage units**.
- 2 On the tape storage unit list, click on the tape storage unit you want to edit.
- 3 Click **Edit** and make the required changes. Click **Save** after making the changes.

Copy a storage unit

You can copy a storage unit to create a new storage unit with the same settings. This option is not available for **OST** storage type.

See [the section called “Copy a disk storage unit”](#) on page 270.

See [the section called “Copy a tape storage unit”](#) on page 271.

Copy a disk storage unit

To copy a disk storage unit

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Storage units**. Select the **Storage units** tab.
- 3 Select the storage unit that you want to copy and select the **Copy storage unit** button.

- 4 Type a unique name for the new storage unit. For example, describe the type of storage. Use this name to specify a storage unit for policies and schedules.

See [“NetBackup naming conventions”](#) on page 614.

- 5 Edit the other properties and disk pool as necessary.
- 6 After reviewing the changes, select the **Save** button.

Copy a tape storage unit

To copy a tape storage unit

- 1 On the left, click on the **Storage > Storage units**.
- 2 Select the tape storage unit that you want to copy and select the **Copy storage unit** button.

Note that the storage unit name is appended with “_Copy”.

- 3 Make any changes that you want and select the **Save** button.

Delete a storage unit

To delete a storage unit from a NetBackup configuration means to delete the label that NetBackup associates with the physical storage.

Deleting a storage unit does not prevent files from being restored that were written to that storage unit. (As long as the storage was not physically removed and the backup image has not expired.)

To delete a storage unit

- 1 Open the NetBackup web UI.
- 2 Use the **Catalog** utility to expire any images that exist on the storage unit. This action removes the image from the NetBackup catalog.

See [“Expire backup images”](#) on page 403.

- Do not manually remove images from storage unit.
- After the images are expired, they cannot be restored unless the images are imported.

See [“About importing backup images”](#) on page 404.

NetBackup automatically deletes any image fragments from a disk storage unit or a disk pool. This deletion generally occurs within seconds of expiring an image. However, to make sure that all of the fragments are deleted, confirm that the directory is empty on the storage unit.

- 3 On the left, select **Storage > Storage units**. Select the **Storage units** tab..

- 4 Select the storage unit that you want to delete.
- 5 Select **Delete > Yes**.
- 6 Modify any policy that uses a deleted storage unit to use another storage unit.
If a storage unit points to a disk pool, you can delete the storage unit without affecting the disk pool.

Tape storage unit considerations

To create a storage unit of a tape robot or a standalone tape drive, when you add a storage unit select **Drives and robots** as the **Storage type**.

When NetBackup sends a job to a **Tape** storage unit, a request is made to mount the volume in a drive.

If a standalone drive does not contain media or if a required volume is not available to a robot, a mount request appears in the **Pending requests**. (In the NetBackup web UI, open **Tape storage > Device monitor**). An operator can then find the volume, mount it manually, and assign it to the drive.

Take the following items into consideration when you add a storage unit for drives and robots:

- Where to add the storage unit depends on which version of NetBackup is in use.
 - Add the storage unit to the primary server. Specify the media server where the drives attach.
 - If using NetBackup Server, add the storage unit to the primary server where the drives attach. The robotic control must also attach to that server.
- The number of storage units that you must create for a robot depends on the robot's drive configuration.
 - Drives with identical densities must share the same storage unit on the same media server. If a robot contains two drives of the same density on the same media server, add only a single storage unit for the robot. Set the **Maximum concurrent write drives** setting to 2.
 - Drives with different densities must be in separate storage units. Consider an STK SL500 library that is configured as a Tape Library DLT (TLD). It can have both half-inch cartridge and DLT drives. Here, you must define a separate storage unit for each density.
 - If a robot's drives and robotic control attach to different NetBackup servers, specify the server where the drives attach as the media server. Always specify the same robot number for the drives as is used for the robotic control.

Disk storage unit considerations

NetBackup permits the creation of an unlimited number of disk storage units.

[Table 13-2](#) describes the different disk types that NetBackup can use as disk media.

Table 13-2 Disk media descriptions

Type of disk storage unit	Description
Media Server Deduplication Pool	A Media Server Deduplication Pool disk type storage unit is used for deduplicated data for a Media Server Deduplication Pool (MSDP), MSDP Cloud, or an MSDP volume group (MVG).
AdvancedDisk	An AdvancedDisk disk type storage unit is used for a dedicated disk that is directly attached to a NetBackup media server. An AdvancedDisk selection is available only when the Data Protection Optimization Option is licensed. NetBackup assumes the exclusive ownership of the disk resources that comprise an AdvancedDisk disk pool. If the resources are shared with other users, NetBackup cannot manage disk pool capacity or storage lifecycle policies correctly. For AdvancedDisk, the NetBackup media servers function as both data movers and storage servers. See the NetBackup AdvancedDisk Storage Solutions Guide.
OpenStorage	An OpenStorage disk type storage unit is used for disk storage, usually provided by a third-party vendor. The actual name of the disk type depends on the vendor. An OpenStorage selection is available only when the OpenStorage Disk Option is licensed. The storage provided by storage vendor partners is integrated into NetBackup via the API. The storage host is the storage server. The NetBackup media servers function as the data movers. The storage vendor's plug-in must be installed on each media server that functions as a data mover. The logon credentials to the storage server must be configured on each media server. See the NetBackup OpenStorage Solutions Guide for Disk.
Cloud Connector	An Cloud Connector disk type storage unit is used for backups to cloud object storage, with no deduplication. See the NetBackup for Cloud Object Store Administrator's Guide.

Table 13-2 Disk media descriptions (*continued*)

Type of disk storage unit	Description
BasicDisk	A BasicDisk type storage unit consists of a directory on a locally-attached disk or a network-attached disk that is exposed as a file system to a NetBackup media server. NetBackup stores backup data in the specified directory. Notes about the BasicDisk type storage unit: ■ Do not include the same volume or file system in multiple BasicDisk storage units. ■ BasicDisk storage units cannot be used in a storage lifecycle policy.

Not all settings are available on each disk storage unit type.

Note: It is recommended that you do not impose quotas on any file systems that NetBackup uses for disk storage units. Some NetBackup features may not work properly when file systems have quotas in place. (For example, the capacity-managed retention selection in storage lifecycle policies and staging to storage units.)

About the disk storage model

The NetBackup model for disk storage accommodates all Enterprise Disk Options. That is, it is the model for all disk types except for the BasicDisk type.

The following items describe components of the disk storage model:

Data mover

An entity that moves data between the primary storage (the NetBackup client) and the storage server. NetBackup media servers function as data movers.

Depending on the disk option, a NetBackup media server also may function as a storage server.

Storage server

An entity that writes data to and reads data from the disk storage. A storage server is the entity that has a mount on the file system on the storage.

Depending on the NetBackup option, the storage server is one of the following:

- A computer that hosts the storage. The computer may be embedded in the storage device.
- A storage vendor's host on the Internet that exposes cloud storage to NetBackup. Alternatively, private cloud storage can be hosted within your private network.

- A NetBackup media server that hosts storage.

Disk pool

A collection of disk volumes that are administered as an entity. NetBackup aggregates the disk volumes into pools of storage (a disk pool) you can use for backups.

A disk pool is a storage type in NetBackup. When you create a storage unit, you select the disk type and then you select a specific disk pool.

Configure the NetBackup service credentials for CIFS storage and disk storage units

For Common Internet File System (CIFS) storage with AdvancedDisk and BasicDisk storage units, the NetBackup Client Service and the NetBackup Remote Manager and Monitor Service services on Windows computers must use the same account credentials. If the account credentials are not configured properly, NetBackup marks all CIFS AdvancedDisk and BasicDisk storage units that use the UNC naming convention as DOWN.

To configure service credentials for CIFS storage and disk storage units

- 1 On the media servers that have a file system mount on the CIFS storage, configure the account and the credentials.

The account must be the same account that the Windows operating system uses for read and write access to the CIFS share.

- 2 In Windows, configure both the NetBackup Client Service and the NetBackup Remote Manager and Monitor Service to run under the same Windows user account that you created in step 1.

See your Windows documentation for details on how to configure the account for the services.

Disk storage units in storage lifecycle policies

Figure 13-1 is an example of how storage lifecycle policies can interact with volumes in a disk pool that a storage unit references.

Two backup policies are created as follows:

- A backup policy named Policy_gold has a gold classification. For storage, it is configured to use an SLP named Lifecycle_Gold, which has a gold data classification.


- A backup policy named Policy_silver has a silver classification. For storage, it is configured to use **Any Available**. That means it can use any available storage unit or any SLP that has a silver classification.

Two storage units are available to both backup policies as follows:

- DSU_1 is an operation in Lifecycle_Gold and references DiskPool_A.
- DSU_2 is not in an SLP and references DiskPool_A.

DiskPool_A contains three disk volumes. Both the gold and the silver images can be written to any disk volume in the pool.

Figure 13-1 Storage lifecycle policies and disk storage units referencing disk pools



Maintain the available disk space on disk storage units

Disk storage units can be managed so that they do not become entirely full and cause backups to fail.

Create space for more images on a disk storage unit in the following ways:

- Add new disk space.

- Set the **High water mark** to a value that best works with the size of backup images in the environment.

Maintain space on basic disk staging storage units in the following ways:

- Increase the frequency of the relocation schedule. Or, add resources so that all images can be copied to a final destination storage unit in a timely manner.

- Run the `nb_updatedssu` script.

Upon NetBackup installation or upgrade, the `nb_updatedssu` script runs. The script deletes the `.ds` files that were used in previous releases as pointers to relocated data. Relocated data is tracked differently in the current release and the `.ds` files are no longer necessary. Under some circumstances, a `.ds` file cannot be deleted upon installation or upgrade. In that case, run the script again:

On Windows: `install_path\NetBackup\bin\goodies\nb_updatedssu`

On UNIX: `/usr/openv/netbackup/bin/goodies/nb_updatedssu`

- Determine the potential free space.
See [“Finding potential free space on a BasicDisk disk staging storage unit”](#) on page 345.
- Monitor disk space by enabling the **Check the capacity of disk storage units** host property.
This General Server host property determines how often NetBackup checks 6.0 disk storage units for available capacity. Subsequent releases use internal methods to monitor disk space more frequently.
See [“General server properties”](#) on page 144.

NDMP storage unit considerations

The NetBackup for NDMP license must be installed on the media server to use the hosts as storage units. Media Manager controls NDMP storage units but the units attach to NDMP hosts.

Create NDMP storage units for drives directly attached to NAS filers. Any drive that is attached to a NetBackup media server is considered a Media Manager storage unit, even if used for NDMP backups.

Note: Remote NDMP storage units may already be configured on a media server from a previous release. Upon upgrade of the media server, those storage units are automatically converted to Media Manager storage units.

See the [NetBackup for NDMP Administrator's Guide](#) for more information.

Configuring robots and tape drives

This chapter includes the following topics:

- [NetBackup robot types](#)
- [Prerequisite for configuring robots and drives](#)
- [About configuring robots and tapes drives in NetBackup](#)
- [Configure drives and robots by using the wizard](#)
- [Configure drive name rules](#)
- [Updating the device configuration by using the wizard](#)
- [Robot properties and configuration options](#)
- [Robot control \(robot configuration options\)](#)
- [Managing robots](#)
- [Managing tape drives](#)

NetBackup robot types

A robot is a peripheral device that moves tape volumes into and out of tape drives. NetBackup uses robotic control software to communicate with the robot firmware.

NetBackup classifies robots according to one or more of the following characteristics:

- The communication method the robotic control software uses; SCSI and API are the two main methods.

- The physical characteristics of the robot. Library refers to a large robot, in terms of slot capacity or number of drives.
- The media type commonly used by that class of robots. HCART (1/2-inch cartridge tape) is an example of a media type.

The table lists the NetBackup robot types that are supported in release 11.0, with drive and slot limits for each type.

To determine which robot type applies to the model of robot that you use, see the [NetBackup Enterprise Server and Server - Hardware and Cloud Storage Compatibility List](#) for your release.

Table 14-1 NetBackup robot types in release 11.0

Robot type	Description	Drive limits	Slot limits	Note
ACS	Automated Cartridge System	1680	No limit	API control. The ACS library software host determines the drive limit.
TLD	Tape library DLT	No limit	32000	SCSI control.

Note: The user interface for NetBackup may show configuration options for the peripheral devices that are not supported in that release. Those devices may be supported in an earlier release, and a NetBackup primary server can manage the hosts that run earlier NetBackup versions. Therefore, the configuration information for such devices must appear in the user interface. The NetBackup documentation may also describe the configuration information for such devices. To determine which versions of NetBackup support which peripheral devices, see the [NetBackup Enterprise Server and Server - Hardware and Cloud Storage Compatibility List](#).

Prerequisite for configuring robots and drives

- **NDMP credentials:** NDMP (Network Data Management Protocol) host credentials are used to authenticate and manage access to NDMP servers. The credentials require NDMP host name, username, and password.
- **ACS credentials:** An ACS (Access Control Service) host is typically used in the context of managing and authorizing access to various services and applications. In NetBackup web UI, you require ACS device host and ACS host.

About configuring robots and tapes drives in NetBackup

Before you configure the robots and tape drives in NetBackup, they must be attached to the computer and recognized by the operating system. The server platforms that NetBackup supports may require operating system configuration changes to allow device discovery.

Device configuration wizard

Cohesity recommends to use the Device Configuration Wizard to add, modify, update, and delete the following types of devices in NetBackup.

- Robots, including those attached to NDMP and ACS hosts.
- Tape drives, including those attached to NDMP and ACS hosts.
- Shared drives (for NetBackup Shared Storage Option configurations only.)

The wizard discovers the devices that are attached to the media servers and helps to configure them.

About drive name rules

The drive name rules define the rules that NetBackup uses to name drives. NetBackup has a default, global drive rule that is used for all connected device hosts. You can also create drive name rules for specific device hosts (each device host can have its own rule). Host-specific rules override the global rule for the devices that are attached to the specified host.

The default, global drive name rule creates names in the following format:

vendor ID.product ID.index

For example: When you use the default global rule to add a drive, the drive names are as follows:

- The first drive name is: <first_drive_name>.000
- The second drive name is: <second_drive_name>.001

A global rule must always exist and only one global rule can exist. If you want to override the existing global rule, you can create another global rule to replace it. If you remove the global drive rule, a NetBackup creates a new global rule with the default settings.

Drive names are limited to 48 characters. Use any of the following drive attributes as part of a drive name rule:

- Host name
- Robot number
- Robot type
- Drive position
- Drive type
- Serial number
- Vendor ID
- Product ID
- Index

Add custom text field is also available which accepts any of the allowable drive name characters.

Configure drives and robots by using the wizard

Note the following when you use the **Configure drives and robots** wizard:

- NetBackup supports only one device configuration scan at a time. If there is an ongoing configuration, a new scan is not initiated. (This limitation applies both if a different user performs the scan or if the same user performs the scan in a different session.) In this case, a warning message is displayed. However, if you are an administrator or the same user who has initiated the process, you can terminate the task.
- **Warning:** If you select **Previous** in the **Configure drives and robots** wizard, the current scan is canceled. If you then select **Next**, a new scan is initiated.

To configure robots and drives using the wizard

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**.
- 3 At the top right, select **Actions > Configure storage devices**.

A list of hosts is displayed.

Note: You can launch the wizard from **Settings > Guided setup**. Select **Add storage**. From the **Storage type** list, select **Drives and robots**.

- 4 Select the device hosts that you want to scan to discover the devices.

The wizard cannot automatically configure any devices that are attached to an NDMP server or an ACS robot.

To perform the scan operation on a specific host for these types of devices:

- Locate the device host in the list and select **Edit**.
- Select the type of device.
- Select **Change**.

5 (Conditional) In the **NDMP** step, configure the NDMP host name and credentials.

Select **Add**.

Specify the required information and then select **Add**.

Select the NDMP host that you want to configure.

Select **Next**.

6 (Conditional) In the **ACS** step, configure an ACS host.

Select **Add**.

Specify the required information and select the ACS host that you want to configure.

Select **Next** to initiate the scanning process.

7 The **Scanned hosts** step displays the scan results and the number of devices that are discovered for each host. NetBackup exits the wizard if there are no removable media device for any the selected device hosts.

8 In the **Backup devices** step, you can perform the following actions:

- Select the robots and drives that you want to configure.
 If you do not want to configure a robot or a drive in a robot, clear the check box for that robot or drive. If you want NetBackup to use a specific robotic drive, you must enable the robot.
- Modify the properties of a drive.
 Locate the robot or drive in the list, then select **Actions > Properties**. You can modify the **Robot type** or the **Drive type**. When you have finished your changes, select **Save**.

9 Select **Next**.

10 To commit the changes to the NetBackup device configuration, select **Commit changes**.

The **Updates** step displays the progress of the updates to device configuration.

- 11 If the devices are already configured in NetBackup, no further action is required. If new devices are discovered, to make the devices available to NetBackup, select **Create storage units**.

If you select **Close**, the new devices that NetBackup discovered have no associated storage units and cannot be used for backups. In this case you must run the **Configure storage devices** wizard again later to create the necessary storage units or manually add the storage units in **Storage > Storage units**.

- 12 Locate the storage unit that you want to configure and select **Edit**.

You can edit the following options: **Storage unit name**, **Maximum concurrent write drives**, **Maximum multiplexing per drive**, **Reduce fragment size to**, and **On demand only**.

- 13 Select **Save** to save the storage units.

If you select **Cancel**, the new devices won't have any associated storage units and cannot be used for backups. In this case you must run the **Configure storage devices** wizard again later to create the necessary storage units or manually add the storage units in **Storage > Storage units**.

Configure drive name rules

Use the following procedure to configure the rules that NetBackup uses to name tape drives. These rules are defined in the **Configure drives and robots** wizard.

To configure drive name rules

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**.
- 3 At the top right, select **Actions > Configure storage devices**.
- 4 On the **Device hosts** step of the wizard, go to the section **Configured drive name rules**.
- 5 To create a new rule, select **Add**.
- 6 You can replace the global rule or create a local rule:
 - To replace the global rule, select the **Global rule** check box.
A global rule must always exist and only one global rule can exist. If you want to override the existing global rule, you can create another global rule to replace it.
 - To create a local rule, select **Add** and enter or select the name of the device host.

- 7 Select **Next**.
 - 8 Configure the rules for naming drives:
 - To select a field that you want to use in the drive name, select **Add field**. Then select the field from the dropdown list.
If you use **host name** in the rule and the drive is a shared drive, the name of the first host that discovers the drive is used as the host name. The name for a shared drive must be identical on all servers that share the drive.
 - To add custom text to the drive name rule, select **Add custom text**. Then enter the custom text.
 - To change the order of the fields that are used for the drive name, select **Actions > Move up** or **Actions > Move down**.
 - 9 Select **Create** to finalize the rule.
- See [“About drive name rules”](#) on page 280.

Updating the device configuration by using the wizard

Cohesity recommends using the configure the robots and drives wizard to update the NetBackup device configuration when hardware changes occur.

Update the configuration for all storage device changes.

For example, if you add or delete a robot or drive or add a new SCSI adapter in a host, update the configuration. Do not update the device configuration during backup or restore activity.

To update the device configuration by using the wizard

1. In the NetBackup web UI, expand or select **Storage** and then select **Tape storage**. A list of configured drives are displayed.
2. Select **Robots** tab.
3. Select kebab menu (three vertical dots) and then select **Edit**.
4. Modify the required configuration values and then select **Save**.

Robot properties and configuration options

This topic describes the robot properties.

Device host

The host to which the device is attached.

Robot type

The type of robot. To change the robot type, you must run the **Configure devices** wizard. To locate the robot type to use for specific vendors and models, See the [NetBackup Enterprise Server and Server - Hardware and Cloud Storage Compatibility List](#).

Robot number

The unique, logical identification number for the robotic library. This number identifies the robotic library in displays (for example, TLD (21)) and is also used when you add media for the robot.

Robot control

The robot control properties that you can configure depend on how the robot is controlled.

See “[Robot control \(robot configuration options\)](#)” on page 285.

Robot control (robot configuration options)

The **Robot control** section of the prompt specifies the type of control for the robot. The options that you configure depend on the robot type and the media server type.

Table 14-2 Robot configuration properties

Property	Description
Robot is controlled locally by this device host	Specifies that the host to which the robot is attached controls the robot. You must configure other options (depending on the robot type and device host type).
Robot is handled by a remote host	Specifies that a host other than the device host controls the robot. You must configure other options (based on the selected robot type and device host platform).
Robot control is attached to an NDMP host	Specifies that an NDMP host controls the robot. You must configure other options (depending on the robot type and device host type).

Table 14-2 Robot configuration properties (*continued*)

Property	Description
ACSLS host	Specifies the name of the Sun StorageTek ACSLS host; the ACS library software resides on the ACSLS host. On some UNIX server platforms, this host can also be a media server. The ACS library software component can be any of the following: ■ Automated Cartridge System Library Software (ACSLS) Examples are available in the NetBackup Device Configuration Guide. ■ STK Library Station ■ Storagenet 6000 Storage Domain Manager (SN6000). This STK hardware serves as a proxy to another ACS library software component (such as ACSLS). Note: If the device host that has drives under ACS robotic control is a Windows server, STK LibAttach software must also be installed. Obtain the appropriate LibAttach software from STK. For compatibility information, see the <i>NetBackup Enterprise Server and Server - Hardware and Cloud Storage Compatibility List</i>: http://www.netbackup.com/compatibility
NDMP host name	Specifies the name of the NDMP host to which the robot is attached.
Robot control host	Specifies the host that controls the robot. The name of the host on which the robot information is defined for TLD robots.
Robot device	The following applies to a Windows device host only. Specifies the name of the robot device. Select Browse and then select a robot from the prompt list displaying the Devices. If the discovery operation fails to discover a robot, select Add manually. Enter either the Port, Bus, Target, and LUN numbers or the device name in the next prompt. If the browse operation fails for any other reason, you are prompted to specify the required information. Use the Windows management tools to find the Port, Bus, Target, and LUN numbers.

Table 14-2 Robot configuration properties (*continued*)

Property	Description
Robotic device file	UNIX device host only. Specifies the device file that is used for SCSI connections. The device files are located in the <code>/dev</code> directory tree on the device host. To specify the robotic device file, select Browse and then select a robotic device file from the list. A list of device files is displayed. Select the appropriate file and select the Select button. If the browse operation fails to show all of the attached robots, click Add manually. Enter the path of the device file in the Robotic device file field and then select the Select button. If the browse operation does not find attached robots, you are prompted to specify the information. Information about how to add device files is available in the NetBackup Device Configuration Guide.
Robot device path	NDMP host only. Specifies the name of the robotic device that is attached to the NDMP host.
Port, Bus, Target, LUN	Windows hosts only. The Port, Bus, Target, and LUN are the SCSI coordinates for the robotic device. To specify the SCSI coordinates of the device, enter the Port, Bus, Target, and LUN.

Managing robots

You can perform various tasks to manage your robots.

See [“Change the robot control properties of a robot”](#) on page 287.

See [“Delete a robot”](#) on page 288.

Change the robot control properties of a robot

Use the following procedure to change the configuration information for a robot.

To change the robot control properties of a robot

- 1 Open the **NetBackup web UI**.
- 2 On the left, click **Storage > Tape storage**. Then click the **Robots** tab.
- 3 Select a robot and click **Edit**.

- 4 Change the robot control properties as necessary.

The properties that you can change depend on the robot type, the host type, and the robot control selection.

See [“Change the robot control properties of a robot”](#) on page 287.

- 5 Click **Save**.

If you restart the Device Manager or the device daemon, any backups, archives, or restores that are in progress also may be stopped.

Delete a robot

Use the following procedure to delete a robot or robots when the media server is up and running.

Any drives that are configured as residing in a robot that you delete are changed to standalone drives.

Any media in the deleted robot is also moved to standalone. If the media is no longer usable or valid, delete it from the NetBackup configuration.

See [“Delete a volume”](#) on page 308.

To delete a robot

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**. Click the **Robots** tab.
- 3 Select the robot or robots you want to delete.
- 4 Click **Delete > Delete**.

Note: It may take a few minutes for the web UI to reflect that the robot is deleted. You are prompted to restart the Media Manager device daemon.

Managing tape drives

You can perform various tasks to manage tape drives.

To manage tape drives, open the NetBackup web UI. Then on the left click **Storage > Tape storage**.

See [“Change a drive comment”](#) on page 289.

See [“About downed drives”](#) on page 289.

See [“Change a drive operating mode”](#) on page 290.

See [“Change the operating mode for a drive path”](#) on page 290.

See [“Clean a tape drive”](#) on page 291.

See [“Delete a drive”](#) on page 291.

See [“Reset a drive”](#) on page 292.

See [“Reset the mount time of a drive”](#) on page 292.

See [“Set the drive cleaning frequency”](#) on page 293.

See [“View drive details”](#) on page 293.

Change a drive comment

You can change the comment that is associated with a drive.

To change a drive comment

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Then select the **Device monitor** tab.
- 3 Select a drive.
- 4 Select **Actions > Change drive comment**.
- 5 Add a comment or change the current drive comment.
See [“NetBackup naming conventions”](#) on page 614.
- 6 Click **Save**.

About downed drives

NetBackup downs a drive automatically when there are read or write errors that surpass the threshold within the time window. The default drive error threshold is 2. That is, NetBackup downs a drive on the third drive error in the default time window (12 hours).

Common reasons for write failures are dirty write heads or old media. The reason for the action is logged in the NetBackup error catalog (view the Media Logs report or the All Log Entries report). If NetBackup downs a device, it is logged in the system log.

You can use the NetBackup `nbemmcmd` command with the `-drive_error_threshold` and `-time_window` options to change the default values.

See [“Change a drive operating mode”](#) on page 290.

Change a drive operating mode

Usually you do not need to change the operating mode of a drive. When you add a drive, NetBackup sets the drive state to UP in Automatic Volume Recognition (AVR) mode. Other operating mode settings are used for special purposes.

The drive operating mode is displayed and changed on the **Device monitor** tab.

To change the mode of a drive

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**. Click the **Device monitor** tab.
- 3 Select a drive or multiple drives.
- 4 Select the operating mode that you want to apply to the drive. For example, to change a drive's status to **Down**, select **Down drive**.

Note that **Up Drive**, **Operator control** applies only to standalone drives.

- 5 If the drive is configured with multiple device paths or is a shared drive (Shared Storage Option), a screen displays that contains a list of all the device paths to the drive. Select the path or paths to change.

Change the operating mode for a drive path

The Device monitor shows path information for drives, including the following:

- Multiple (redundant) paths to a drive are configured
- Any drives are configured as shared drives (Shared Storage Option)

To change the operating mode for a drive path

- 1 Open the web UI.
- 2 On the left, click **Storage > Tape storage**. Click the **Device monitor** tab.
- 3 Click on the drive name to view the drive properties. Then click on the **Paths** tab.
- 4 Select a path or select multiple paths.
- 5 Click **Actions**, then choose a command for the path action, as follows:
 - **Up path**
 - **Down path**
 - **Reset path**

Clean a tape drive

When you add a drive to NetBackup, you can configure the automatic, frequency-based cleaning interval.

You can also perform an operator-initiated cleaning of a drive regardless of the cleaning frequency or accumulated mount time of the drive. However, appropriate cleaning media must be added to NetBackup.

After you clean a drive, reset the mount time.

See [“Reset the mount time of a drive ”](#) on page 292.

To clean a tape drive

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Select the **Device monitor** tab.
- 3 Select the drive to clean.
- 4 Select **Actions > Drive cleaning > Clean now**. NetBackup initiates drive cleaning regardless of the cleaning frequency or accumulated mount time.

The **Clean now** option resets the mount time to zero, but the cleaning frequency value remains the same. If the drive is a standalone drive and it contains a cleaning tape, NetBackup issues a mount request.

- 5 For a shared drive (Shared Storage Option), do the following:

In the list of hosts that share the drive, choose only one host on which the function applies.

- 6 Select **Clean now**.

The **Clean now** function can take several minutes to complete, so the cleaning information may not update immediately.

Delete a drive

Use the following procedure to delete a drive or drives when the media server is up and running.

To delete a drive

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Select the **Device monitor** tab.

- 3 Select the required drive.
- 4 Select **Delete**.

Note: It may take a few minutes for the web UI to reflect that the drive is deleted. You are prompted to restart the Media Manager device daemon.

Reset a drive

Resetting a drive changes the state of the drive.

Usually you reset a drive when its state is unknown, which occurs if an application other than NetBackup uses the drive. When you reset the drive, it returns to a known state before use with NetBackup. If a SCSI reservation exists on the drive, a reset operation from the host that owns the reservation can help the SCSI reservation.

If the drive is in use by NetBackup, the reset action fails. If the drive is not in use by NetBackup, NetBackup tries to unload the drive and set its run-time attributes to default values.

Note that a drive reset does not perform any SCSI bus or SCSI device resets.

To reset a drive

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Select the **Device monitor** tab.
- 3 Locate the drive that you want to reset. Then select **Actions > Reset drive**.
- 4 If the drive is in use by NetBackup and cannot be reset, restart the NetBackup Job Manager (`nbjmgr`) to free up the drive.
- 5 Determine which job controls the drive (that is, which job writes to or reads from the drive).

On the left, select **Activity monitor**. Then on the **Jobs** tab, cancel the job.

- 6 In the **Activity monitor**, restart the NetBackup Job Manager, which cancels all NetBackup jobs in progress.

Reset the mount time of a drive

You can reset the mount time of the drive. Reset the mount time to zero after you perform a manual cleaning.

To reset the mount time

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**. Click the **Device monitor** tab.
- 3 Select a drive.

- 4 Click **Actions > Drive cleaning > Reset mount time**. The mount time for the selected drive is set to zero.
- 5 If you use the Shared drive (Shared Storage Option), do the following:
In the list of hosts that share the drive, choose only one host on which the function applies.
- 6 Click **Reset mount time**.

Set the drive cleaning frequency

When you add a drive to NetBackup, you configure the automatic, frequency-based cleaning interval. From the **Device monitor** you can change the cleaning frequency that was configured when you added the drive.

To set the cleaning frequency

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**. Click the **Device monitor** tab.
- 3 Select a drive.
- 4 Click **Actions > Drive cleaning > Set cleaning frequency**.
- 5 Enter the number of mount hours between drive cleaning.

The **Set cleaning frequency** option is not available for the drives that do not support frequency-based cleaning. This function is not available for shared drives.

The drive cleaning interval appears in the **Drive properties**.
- 6 Click **Save**.

View drive details

You can obtain detailed information about drives (or shared drives), such as drive cleaning, drive properties, drive status, host, and robotic library information.

To view the drive details

- 1 Open the web UI.
- 2 On the left, select **Storage > Tape storage**. Select the **Device monitor** tab.
- 3 Many drive details are displayed on this tab. For additional details, select the link for a drive name.

For shared drives, you can see the drive **Control** mode and **Drive index** for each host that shares a drive. Click on the **Shared drive hosts** tab to view a list of hosts that share a drive.

Configuring tape media

This chapter includes the following topics:

- [About NetBackup tape volumes](#)
- [About NetBackup volume pools](#)
- [About NetBackup volume groups](#)
- [NetBackup media types](#)
- [About adding volumes](#)
- [Managing volumes](#)
- [Managing volume pools](#)
- [Managing volume groups](#)

About NetBackup tape volumes

A tape volume is a data storage tape or a cleaning tape. NetBackup assigns attributes to each volume and uses them to track and manage the volumes. Attributes include the media ID, robot host, robot type, robot number, and slot location.

NetBackup uses two volume types, as follows:

Robotic volumes

Volumes that are located in a robot.

The robotic library moves the volumes into and out from the robotic drives as necessary.

Standalone volumes	Volumes that are allocated for the drives that are not in a robot. Operator intervention is required to load volumes into and eject volumes from standalone drives.

NetBackup uses volume pools to organized volumes by usage.

See “[About NetBackup volume pools](#)” on page 295.

Volume information is stored in the NetBackup database.

About NetBackup volume pools

A volume pool identifies a set of volumes by usage. Volume pools protect volumes from access by unauthorized users, groups, or applications. When you add media to NetBackup, you assign them to a volume pool (or assign them as standalone volumes, without a pool assignment).

By default, NetBackup creates the following volume pools:

NetBackup	The default pool to which all backup images are written (unless you specify otherwise).
DataStore	For DataStore use.
CatalogBackup	For NetBackup catalog backups. Catalog backup volumes are not a special type in NetBackup. They are the data storage volumes that you assign to the CatalogBackup volume pool. To add NetBackup catalog backups, use any of the add volume methods. Ensure that you assign them to the volume pool you use for catalog backups. After adding volumes, use the NetBackup Catalog Backup wizard to configure a catalog backup policy.
None	For the volumes that are not assigned to a pool.

You can add other volume pools. For example, you can add a volume pool for each storage application you use. Then, as you add volumes to use with an application, you assign them to that application’s volume pool. You can also move volumes between pools.

You also can configure a scratch pool from which NetBackup can transfer volumes when a volume pool has no volumes available.

The volume pool concept is relevant only for tape storage units and does not apply to disk storage units.

You can use any of the approved characters for volume pool names.

NetBackup uses several special prefixes for volume pool names.

About NetBackup volume groups

A volume group identifies a set of volumes that reside at the same physical location. The location can be either the robot in which the volumes reside, standalone storage, or off-site storage if you use the NetBackup Vault option.

When you add media to NetBackup, NetBackup assigns all volumes in a robot to that robot's volume group. Alternatively, you can assign the media to a different group.

Volume groups are convenient for tracking the location of volumes, such as the case when a volume is moved off site. Volume groups let you perform operations on a set of volumes by specifying the group name rather than each individual media ID of each volume. Operations include moves between a robotic library and a standalone location or deletions from NetBackup.

If you move a volume physically, you also must move it logically. A logical move means to change the volume attributes to show the new location.

The following are the rules for assigning volume groups:

- All volumes in a group must be the same media type.
However, a media type and its corresponding cleaning media type are allowed in the same volume group (such as DLT and DLT_CLN).
- All volumes in a robotic library must belong to a volume group.
You cannot add volumes to a robotic library without specifying a group or having Media Manager generate a name for the group.
- The only way to clear a volume group name is to move the volume to standalone and not specify a volume group.
- More than one volume group can share the same location.
For example, a robotic library can contain volumes from more than one volume group and you can have more than one standalone volume group.
- All volumes in a group must be in the same robotic library or be standalone.
That is, you cannot add a group (or part of a group) to a robotic library if it already exists in another robotic library.

Examples of volume group usage are available.

NetBackup media types

NetBackup uses media types to differentiate the media that have different physical characteristics. Each media type may represent a specific physical media type.

The NetBackup media types are also known as Media Manager media types.

The following table describes the NetBackup media types.

Table 15-1 NetBackup media types

Media type	Description
DLT	DLT cartridge tape
DLT_CLN	DLT cleaning tape
DLT2	DLT cartridge tape 2
DLT2_CLN	DLT cleaning tape 2
DLT3	DLT cartridge tape 3
DLT3_CLN	DLT cleaning tape 3
HCART	1/2 inch cartridge tape
HCART2	1/2 inch cartridge tape 2
HCART3	1/2 inch cartridge tape 3
HC_CLN	1/2 inch cleaning tape
HC2_CLN	1/2 inch cleaning tape 2
HC3_CLN	1/2 inch cleaning tape 3

NetBackup writes media in a format that allows the position to be verified before NetBackup appends new backup images to the media.

Note: The user interface for NetBackup may show configuration options for the media types that are not supported in that release. Those types may be supported in an earlier release, and a NetBackup primary server can manage the hosts that run earlier NetBackup versions. Therefore, the configuration information for such types must appear in the user interface. The NetBackup documentation also may describe the configuration information for such types. To determine which versions of NetBackup support which media types, see the *NetBackup Enterprise Server and Server - Hardware and Cloud Storage Compatibility List*:

<http://www.netbackup.com/compatibility>

Alternate NetBackup media types

Alternate media types let you define more than one type of tape in the same library. You can use the alternate types to differentiate between different physical cartridges.

The following are examples of alternate media types:

- DLT, DLT2, DLT3
- HCART, HCART2, HCART3

For example, if a robot has DLT4000 and DLT7000 drives, you can specify the following media types:

- DLT media type for the DLT4000 tapes
- DLT2 media type for the DLT7000 tapes

NetBackup then does not load a tape that was written in a DLT4000 drive into a DLT7000 drive and vice versa.

You must use the appropriate default media type when you configure the drives. (When you configure drives in NetBackup, you specify the default media type to use in each drive type.)

In a robot, all of the volumes (of a specific vendor media type) must be the same NetBackup media type. For example, for an ACS robot that contains 3490E media, you can assign either NetBackup HCART, HCART2, or HCART3 media type to that media. You cannot assign HCART to some of the media and HCART2 (or HCART3) to other of the media.

About adding volumes

Adding volumes is a logical operation that assigns NetBackup attributes to physical media. The media can reside in storage devices already, or you can add them to the storage devices when you add them to NetBackup. How you add volumes depends on the type of volume: robotic or standalone.

See [“About adding robotic volumes”](#) on page 299.

See [“About adding standalone volumes”](#) on page 299.

NetBackup uses the rules to assign names and attributes to volumes.

About adding robotic volumes

The robotic volumes are the volumes that are located in a robotic tape library. The following table describes the methods for adding robotic volumes.

Table 15-2 Methods for adding robotic volumes

Method	Description
Manually add a volume	See “Add a volume” on page 300.
Robot inventory	See “About showing a robot's contents” on page 324. See “About comparing a robot's contents with the volume configuration” on page 326. See “About previewing volume configuration changes” on page 327. See “Update the NetBackup volume configuration with a robot's contents” on page 329.
NetBackup commands	See the NetBackup Commands Reference Guide .

About adding standalone volumes

Standalone volumes are the volumes that reside in the drives that are not in a robot or are allocated for standalone drives.

Because NetBackup does not label volumes until it uses them, you can add volumes even though they do not reside in a drive. The additional volumes are available for use if the volume in a drive becomes full or unusable. For example, if a volume in a standalone drive is full or unusable because of errors, NetBackup ejects (logically) the volume. If you add other standalone volumes, NetBackup requests that volume; NetBackup does not generate an `out of media` error.

The `DISABLE_STANDALONE_DRIVE_EXTENSIONS` option of the `nbemmcmd` command can turn off the automatic use of standalone volumes.

Table 15-3 Methods for adding standalone volumes

Method	Description
Manually add a volume	See “Add a volume” on page 300.
NetBackup commands	See the NetBackup Commands Reference Guide .

Add a volume

Use this procedure to add a new volume.

Be careful when you specify properties. You cannot change some properties later, such as the media ID or type. If you specify them incorrectly, you must delete the volume and add it again.

To add a volume

- 1 Open the **NetBackup web UI**.
- 2 On the left, select **Storage > Tape storage**.
- 3 Select the **Volumes** tab.
- 4 Select **Add volume**.
- 5 Specify the properties for the volumes.

The properties that display can vary depending on the type of volume.

See [“Volume properties”](#) on page 300.

- 6 Select **Save**.

If the robot has a barcode reader, NetBackup performs the following actions:

- Adds the volume to the EMM database using the specified media ID.
- Reads the barcode of each new volume.
- Adds the barcodes as attributes in the EMM database.

Volume properties

Volume properties describes the properties for volumes in NetBackup. The properties depend on whether you add, change, or move volumes.

The properties are arranged alphabetically.

Table 15-4 Volume properties

Property	Description	Operation
Device host	The name of the NetBackup media server to which the robot is attached.	Add, move
Expiration date	The following does not apply to cleaning tapes. The date after which the volume is too old to be reliable. When the expiration date has passed, NetBackup reads data on the volume but does not mount and write to the volume. You should exchange it for a new volume. When you add a new volume, NetBackup does not set an expiration date. The expiration date is not the same as the retention period for the backup data on the volume. You specify data retention periods in the backup policies.	Change
First media ID	This property appears only if the number of volumes is more than one. The ID of the first volume in the range of volumes. Media IDs need to be exactly 6 characters. Valid only when you add a range of volumes. Use the same pattern that you chose in the Media ID naming style box. NetBackup uses the pattern to name the remaining volumes by incrementing the digits. NetBackup allows specific characters in names.	Add
First slot number	The number of the first slot in the robot in which the range of volumes resides. If you add or move more than one media, NetBackup assigns the remainder of the slot numbers sequentially. Note: You cannot enter slot information for volumes in an API robot. The robot vendor tracks the slot locations for API robot types.	Add, move
Maximum cleanings	The maximum number of times NetBackup should mount the volume or use the cleaning tape. To determine the maximum mount limit to use, consult the vendor documentation for information on the expected life of the volume.	Add
Maximum mounts	The following topic does not apply to cleaning tapes. The Maximum mounts property specifies the number of times that the selected volumes can be mounted. When the limit is reached, NetBackup reads data on the volume but does not mount and write to the volume. A value of zero (the default) is the same as Unlimited. To help determine the maximum mount limit, consult the vendor documentation for information on the expected life of the volume.	Add, change

Table 15-4 Volume properties (*continued*)

Property	Description	Operation
Media description	A description of the media, up to 25 character maximum. NetBackup allows specific characters in names.	Add, change
Media ID	This property appears only if the number of volumes is one. The ID for the new volume. Media IDs must be exactly 6 characters. Media IDs for an API robot must match the barcode on the media (for API robots, NetBackup supports barcodes of 6 characters). Therefore, obtain a list of the barcodes before you add the volumes. Obtain this information through a robotic inventory or from the robot vendor's software. NetBackup allows specific characters in names.	Add, change
Media ID naming style	The style to use to name the range of volumes. Media IDs must be exactly 6 characters in length. Using the pattern, NetBackup names the remaining volumes by incrementing the digits. NetBackup media IDs for an API robot must match the barcode on the media. For API robots, NetBackup supports barcodes from 1 to 6 characters. Therefore, obtain a list of the barcodes before you add the volumes. Obtain this information through a robotic inventory or from the robot vendor's software. NetBackup allows specific characters in names.	Add
Media type	The media type for the volume to add. Select the type from the drop-down list.	Add
Number of volumes	The number of volumes to add. For a robotic library, enough slots must exist for the volumes.	Add
Robot	The robotic library to add or move the volumes to. To add volumes for a different robot, select a robot from the drop-down list. The list shows robots on the selected host that can contain volumes of the selected media type.	Add, move

Table 15-4 Volume properties (*continued*)

Property	Description	Operation
Volume group	If you specified a robot, select from a volume group already configured for that robot. Alternatively, enter the name for a volume group; if it does not exist, NetBackup creates it and adds the volume to it. If you do not specify a volume group (you leave the volume group blank), the following occurs: ■ Standalone volumes are not assigned to a volume group. ■ NetBackup generates a name for robotic volumes by using the robot number and type. For example, if the robot is a TLD and has a robot number of 50, the group name is 000_00050_TLD. See “About NetBackup volume groups” on page 296. See “About rules for moving volumes between groups” on page 310.	Add, move
Volume is in a robotic library	When you add a volume: ■ If the volume is in a robot, select Volume is in a robotic library. ■ If the volume is a standalone volume, do not select Volume is in a robotic library. When you move a volume: ■ To inject a volume into a robotic library, select Volume is in a robotic library. Then, select a robot and the slot number (First slot number) for the volume. ■ To eject a volume from a robot, clear Volume is in a robotic library.	Add, move
Volume pool	The pool to which the volume or volumes should be assigned. Select a volume pool you created or one of the following standard NetBackup pools: ■ None. ■ NetBackup is the default pool name for NetBackup. ■ DataStore is the default pool name for DataStore. ■ CatalogBackup is the default pool name used for NetBackup catalog backups of policy type NBU-Catalog. When the images on a volume expire, NetBackup returns it to the scratch volume pool if it was allocated from the scratch pool. See “About NetBackup volume pools” on page 295.	Add, change
Volumes to move	The Volumes to move section of the dialog box shows the media IDs of the volumes that you selected to move.	Move

Managing volumes

The following sections describe the procedures to manage volumes.

See [“Edit a volume”](#) on page 304.

See [“Move volumes”](#) on page 305.

See [“About recycling a volume”](#) on page 306.

See [“Delete a volume”](#) on page 308.

See [“Changing the media owner of a volume”](#) on page 309.

See [“Changing the volume group assignment”](#) on page 309.

See [“Rescan and update barcodes”](#) on page 310.

See [“Eject volumes”](#) on page 312.

See [“Label a volume”](#) on page 313.

See [“Erase a volume”](#) on page 314.

See [“Freeze or unfreeze a volume”](#) on page 315.

See [“Suspend or unsuspend volumes”](#) on page 316.

Edit a volume

You can change some of the properties of a volume, including the volume pool.

To change volume properties

- 1 Open the **NetBackup web UI**.
- 2 On the left, select **Storage > Tape storage**.
- 3 Select the **Volumes** tab.
- 4 Locate and select the volume. Select **Edit**.
- 5 Change the properties for the volume.
See [“Volume properties”](#) on page 300.
- 6 Select **Update**.

About moving volumes

When you move volumes in or out of a robotic library or from one robot to another, move the volumes physically and logically, as follows:

- Physically move volumes by inserting or by removing them. For some robot types, use the NetBackup inject and eject options.

- Logically move volumes using NetBackup, which updates the NetBackup database to show the volume at the new location.

When you move volumes from one robotic library to another robotic library, perform the following actions:

- Move the volumes to stand alone as an intermediate step.
- Move the volumes to the new robotic library.

The following types of logical moves are available:

- Move single volumes.
- Move multiple volumes.
- Move combinations of single and multiple volumes.
- Move volume groups.

You cannot move volumes to an invalid location.

It is recommended that you perform moves by selecting and by moving only one type of media at a time to a single destination.

The following are several examples of when to move volumes logically:

- When a volume is full in a robotic library and no slots are available for new volumes in the robotic library. Move the full volume to stand alone, remove it from the robot, then configure a new volume for the empty slot or move an existing volume into that slot. Use the same process to replace a defective volume.
- Moving volumes from a robotic library to an off-site location or from an off-site location into a robotic library. When you move tapes to an off-site location, move them to stand alone.
- Moving volumes from one robotic library to another (for example, if a library is down).
- Changing the volume group for a volume or volumes.

Move volumes

If you move a volume to a robotic library that has a barcode reader, NetBackup updates the EMM database with the correct barcode.

- 1 Open the **NetBackup web UI**.
- 2 On the left, select **Storage > Tape storage**.
- 3 Select the **Volumes** tab.
- 4 Select the desired volumes from the list and select **Move**.

5 Specify the properties for the move.

If you move a single volume, the dialog box entries show the current location of the volume.

See [“Volume properties”](#) on page 300.

6 Select **Confirm**.

About recycling a volume

If you recycle a volume, you can use either the existing media ID or a new media ID.

Caution: Recycle a volume only if all NetBackup data on the volume is no longer needed or if the volume is damaged and unusable. Otherwise, you may encounter serious operational problems and a possible loss of data.

Recycling a volume and using the existing media ID

NetBackup recycles a volume and returns it to the volume rotation when the last valid image on the volume expires.

To recycle a volume that contains unexpired backup images, you must deassign the volume.

See [“About assigning and deassigning volumes”](#) on page 307.

Recycling a volume and using a new media ID

Recycle a volume if it is a duplicate of another volume with the same media ID. Also recycle a volume if you change how you name volumes and you want to match the barcodes on the volume.

The following table describes the procedure to recycle a volume and use a new media ID.

Table 15-5 Recycling a volume and using a new media ID

Step	Action	Description
Step 1	Physically remove the volume from the storage device.	See “Eject volumes” on page 312.
Step 2	If the volume is in a robotic library, move it to standalone.	See “About moving volumes” on page 304.

Table 15-5 Recycling a volume and using a new media ID (*continued*)

Step	Action	Description
Step 3	Record the current number of mounts and expiration date for the volume.	Go to Storage > Tape storage > Volumes in the NetBackup web UI).
Step 4	Delete the volume entry.	See “Delete a volume” on page 308.
Step 5	Add a new volume entry.	See “Add a volume” on page 300. Because NetBackup sets the mount value to zero for new volume entries, you must adjust the value to account for previous mounts. Set the maximum mounts to a value that is equal to or less than the following value: The number of mounts that the manufacturer recommends minus the value that you recorded earlier.
Step 6	Physically add the volume to the storage device.	See “Inject volumes into robots” on page 311.
Step 7	Configure the number of mounts.	Set the number of mounts to the value you recorded earlier by using the following command: On Windows hosts: <pre>install_path\Volmgr\bin\vmchange -m media_id -n number_of_mounts</pre> On UNIX hosts: <pre>/usr/opensv/volmgr/bin/vmchange -m media_id -n number_of_mounts</pre>
Step 8	Set the expiration date to the value you recorded earlier.	See “Edit a volume” on page 304.

About assigning and deassigning volumes

An assigned volume is one that is reserved for exclusive use by NetBackup. A volume is set to the assigned state when either application writes data on it for the first time. The time of the assignment appears in the **Time assigned** column for the volume on the **Volumes** tab. When a volume is assigned, you cannot delete it or change its volume pool.

A volume remains assigned until NetBackup deassigns it.

NetBackup deassigns a volume only when the data is no longer required, as follows:

- For regular backup volumes, when the retention period has expired for all the backups on the volume.
- For catalog backup volumes, when you stop using the volume for catalog backups.

To deassign a volume, you expire the images on the volume. After you expire a volume, NetBackup deassigns it and does not track the backups that are on it. NetBackup can reuse the volume, you can delete it, or you can change its volume pool.

See [“Expire backup images”](#) on page 403.

You can expire backup images regardless of the volume state (Frozen, Suspended, and so on).

NetBackup does not erase images on expired volumes. You can still use the data on the volume by importing the images into NetBackup (if the volume has not been overwritten).

See [“About importing backup images”](#) on page 404.

Note: It is not recommended that you deassign NetBackup volumes. If you do, be certain that the volumes do not contain any important data. If you are uncertain, copy the images to another volume before you deassign the volume.

Delete a volume

You can delete volumes from the NetBackup configuration. For example, if any of the following situations apply, you may want to delete the volume:

- A volume is no longer used and you want to recycle it by relabeling it with a different media ID.
- A volume is unusable because of repeated media errors.
- A volume is past its expiration date or has too many mounts, and you want to replace it with a new volume.
- A volume is lost and you want to remove it from the NetBackup database.

After a volume is deleted, you can discard it or add it back under the same or a different media ID.

See [“About assigning and deassigning volumes”](#) on page 307.

To delete volumes

- 1 Before you delete and reuse or discard a volume, ensure that it does not have any important data. You cannot delete NetBackup volumes if they are assigned.
- 2 Open the **NetBackup web UI**.
- 3 Click **Storage > Tape storage**.
- 4 Click the **Volumes** tab.
- 5 Select the desired volume from the volumes list and click **Delete > Delete**.
- 6 Remove the deleted volume or volumes from the storage device.

Changing the media owner of a volume

You can change the media server or server group that owns the volume.

See [“About NetBackup server groups”](#) on page 263.

To change the owner of a volume

- 1 Open the **NetBackup web UI**.
- 2 On the left, select **Storage > Tape storage**.
- 3 Select the **Volumes** tab.
- 4 Select the volume that you want to change.
- 5 Select **Change media owner**.
- 6 From the **Media server** list, select a media owner.
Only the volumes that belong to a server group display in the list.
- 7 Select **Confirm**.

Changing the volume group assignment

If you move a volume physically to a different robot, change the group of the volume to reflect the move.

See [“About rules for moving volumes between groups”](#) on page 310.

To change the group of a volume

- 1 Open the **NetBackup web UI**.
- 2 On the left, select **Storage > Tape storage**.
- 3 Select the **Volumes** tab.
- 4 Select the volumes that you want to change the volume group assignment for.

- 5 Select **Change volume group**.
- 6 For the **Volume group**, enter the name of the new volume group. Or, select a name from the list.
- 7 Select **Confirm**.

The name change is reflected in the volume list entry for the selected volumes. If you specified a new volume group (which creates a new volume group), the group appears under **Volume groups**.

About rules for moving volumes between groups

The following are the rules for moving volumes between groups:

- The target volume group must contain the same type of media as the source volume group. If the target volume group is empty: The successive volumes that you add to it must match the type of media that you first add to it.
- All volumes in a robotic library must belong to a volume group. If you do not specify a group, NetBackup generates a new volume group name by using the robot number and type.
- More than one volume group can share the same location. For example, a robotic library can contain volumes from more than one volume group and you can have more than one standalone volume group.
- All members of a group must be in the same robotic library or be standalone. That is, if volume group already exists in another robotic library, you cannot add it (or part of it) to a robotic library.

See [“About NetBackup volume groups”](#) on page 296.

See [“About moving volumes”](#) on page 304.

Rescan and update barcodes

Use the following procedure to rescan the media in a robot and to update NetBackup with the barcodes.

Note: Rescan and update barcodes does not apply to volumes in API robot types.

To rescan and update barcodes

- 1 Open the **NetBackup web UI**.
- 2 On the left, select **Storage > Tape storage**.
- 3 Select the **Volumes** tab.

- 4 Select the volumes you want to rescan and update.
- 5 Select **Rescan/update barcodes**.
- 6 Select **Start**.
- 7 The results of the update display in the **Results** section.

About barcode rules

A barcode rule specifies criteria for assigning attributes to new robotic volumes. NetBackup assigns these attributes by using the barcode for the volume that the robotic library provides and your barcode rules.

In NetBackup, you choose whether to use barcode rules when you set up the robot inventory update operation. The barcode rules are stored on the primary server.

Note: NetBackup does not use barcode rules if a volume already uses a barcode.

About injecting and ejecting volumes

Media access port (MAP) functionality differs between robotic libraries. For many libraries, NetBackup opens and closes the MAP as needed. However, some libraries have the front-panel inject and the eject functions that conflict with NetBackup's use of the media access port. And for other libraries, NetBackup requires front-panel interaction by an operator to use the media access port.

Read the operator manual for the library to understand the media access port functionality. Some libraries may not be fully compatible with the inject and eject features of NetBackup unless properly handled. Other libraries may not be compatible at all.

Inject volumes into robots

You can inject volumes into the robots that contain media access ports.

Any volumes to be injected must be in the media access port before the operation begins. If no volumes are in the port, you are not prompted to place volumes in the media access port and the update operation continues.

Each volume in the MAP is moved into the robotic library. If the MAP contains multiple volumes, they are moved to empty slots in the robotic library until the media access port is empty or all the slots are full.

After the volume or volumes are moved, NetBackup updates the volume configuration.

Some robots report only that media access ports are possible. Therefore, the **Empty media access port prior to update** option may be available for some robots that do not contain media access ports.

To inject volumes into the robots that contain media access ports

- 1 Load the volumes in the MAP.
- 2 Inventory the robot.

See [“Update the NetBackup volume configuration with a robot's contents”](#) on page 329.

- 3 Select **Empty media access port prior to update**.

Eject volumes

You can eject single or multiple volumes.

You cannot eject multiple volumes with one operation if they reside in multiple robots.

Operator intervention is only required if the robotic library does not contain a media access port large enough to eject all of the selected volumes. For these robot types, NetBackup prompts an operator to remove the media from the media access port so the eject operation can continue.

See [“Media ejection timeout periods”](#) on page 313.

To eject volumes

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Then select the **Volumes** tab.
- 3 Select one or more volumes that you want to eject.
- 4 Click **Eject from robot**.
- 5 Do one of the following actions:

ACS robots	Select the media access port to use for the ejection, then click Eject .

TLD robots	Click Eject .

The robotic library may not contain a media access port large enough to eject all of the selected volumes. For most robot types, you are prompted to remove the media from the media access port so the eject can continue with the remaining volumes.

See [“NetBackup robot types”](#) on page 278.

Media ejection timeout periods

The media ejection period (the amount of time before an error condition occurs) varies depending on the capability of each robot.

The following table shows the ejection timeout periods for robots.

Table 15-6 Media ejection timeout periods

Robot types	Timeout period
Automated Cartridge System (ACS)	One week
Tape Library DLT (TLD)	30 minutes.

Note: If the media is not removed and a timeout condition occurs, the media is returned to (injected into) the robot. Inventory the robot and eject the media that was returned to the robot.

Some robots do not contain media access ports. For these robots, the operator must remove the volumes from the robot manually.

Note: After you add or remove media manually, use NetBackup to inventory the robot.

Label a volume

If a volume contains valid NetBackup images, deassign the volume so that it can be labeled.

If you want to label media and assign specific media IDs (rather than allow NetBackup to assign IDs), use the `bplabel` command.

Note: If you label a volume, NetBackup cannot restore or import the data that was on the media after you label it.

To label a volume

- 1 Open the **NetBackup web UI**.
- 2 On the left, select **Storage > Tape storage**.
- 3 Select the **Volumes** tab.
- 4 Select a volume that you want to label.
- 5 Select **Label**.

6 Specify the following properties for the label operation.

Media server	Enter name of the media server that controls the drive to write the label.
Verify media label before performing action	Select this option to verify that the media in the drive is the expected media. To overwrite any existing labels on the media, do not select Verify media label before performing action.

7 Select **Confirm**.

Erase a volume

You can erase the data on a volume if the following are true:

- The volume is not assigned.
- The volume contains no valid NetBackup images.
- A single volume is selected to erase.

After NetBackup erases the media, NetBackup writes a label on the media.

If you erase the media, NetBackup cannot restore or import the data on the media.

Note: NetBackup does not support erase functions on NDMP drives.

The following table describes the types of erase.

Table 15-7 Types of erase

Type of erase	Description
Long erase	Rewinds the media and the data is overwritten with a known data pattern. A SCSI long erase is also called a secure erase because it erases the recorded data completely. Note: A long erase is a time-consuming operation and can take as long as 2 hours to 3 hours. For example, it takes about 45 minutes to erase a 4-mm tape on a standalone drive.

Table 15-7 Types of erase (*continued*)

Type of erase	Description
Quick erase	Rewinds the media and an erase gap is recorded on the media. The format of this gap is drive dependent. It can be an end-of-data (EOD) mark or a recorded pattern that the drive does not recognize as data. Some drives do not support a quick erase (such as QUANTUM DLT7000). For the drives that do not support a quick erase, the new tape header that is written acts as an application-specific quick erase.

To erase a volume

- 1 If a volume contains valid NetBackup images, deassign the volume so NetBackup can label it.
- 2 Open the **NetBackup web UI**.
- 3 On the left, select **Storage > Tape storage**.
- 4 Select the **Volumes** tab.
- 5 Select a volume that you want to erase.
- 6 Select **Quick erase** or **Long erase**.
- 7 Specify the name of the media server to initiate the erase operation.

To overwrite any existing labels on the media, do not select **Verify media label before performing action**.
- 8 Select **Confirm** if you are certain you want to start the erase action.

If you selected **Verify media label before performing operation** and the actual volume label does not match the expected label, the media is not erased.

Freeze or unfreeze a volume

NetBackup freezes volumes under certain circumstances. Use the following procedure to manually freeze or unfreeze a volume.

To freeze or unfreeze a volume

- 1 Open the **NetBackup web UI**.
- 2 On the left, select **Storage > Tape storage**.
- 3 Select the **Volumes** tab.
- 4 Select the volume that you want to freeze or unfreeze.
- 5 Select **Freeze** or **Unfreeze**.
- 6 Select **Confirm**.

Suspend or unsuspend volumes

You cannot use a suspended volume for backups until retention periods for all backups on it have expired. At that time, NetBackup deletes the suspended volume from the NetBackup media catalog and unassigns it from NetBackup.

A suspended volume is available for restores. If the backups have expired, first import the backups.

To suspend or unsuspend media

- 1 Open the **NetBackup web UI**.
- 2 On the left, select **Storage > Tape storage**.
- 3 Select the **Volumes** tab.
- 4 Select the volumes that you want to suspend or unsuspend.
- 5 Select **Suspend** or **Unsuspend**.
- 6 Select **Confirm**.

Managing volume pools

The following sections describe the operations you can perform to manage volume pools.

See [“Add a volume pool”](#) on page 316.

See [“Edit or delete a volume pool”](#) on page 317.

Add a volume pool

Use this procedure to add a new volume pool.

To add a volume pool

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Then select the **Volume pools** tab.
- 3 Select the **Add volume pool** button.
- 4 Specify the properties for the volume pool.
See [“Volume pool properties”](#) on page 318.
- 5 You can add volumes to the pool in the following ways:
 - Add new volumes to NetBackup.
 - Change the pool of existing volumes.

Edit or delete a volume pool

Edit a volume pool

Use this procedure to change the properties of a volume pool. The properties you can change include the pool type (scratch pool or catalog backup pool).

To edit a volume pool

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Then select the **Volume pools** tab.
- 3 Select a pool in the **Volume pools** list.
- 4 Select **Edit**.
- 5 Change the attributes for the volume pool.
See [“Volume pool properties”](#) on page 318.

Delete a volume pool

You cannot delete any of the following pools:

- A volume pool that contains volumes
- The **NetBackup** volume pool
- The **None** volume pool
- The default **CatalogBackup** volume pool
- The **DataStore** volume pool

To delete a volume pool

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Then select the **Volume pools** tab.
- 3 Locate the volume pool that you want to delete. Ensure that the volume pool is empty. If the pool is not empty, change the pool name for any volumes in the pool.
- 4 Select the volume pool.
- 5 Select **Delete**.
- 6 Select **Confirm**.

Volume pool properties

You can specify various properties for a volume pool.

Table 15-8 Volume pool properties

Property	Description
Catalog backup pool	Select this option to use this volume pool for catalog backups. This check box creates a dedicated catalog backup pool to be used for NBU-Catalog policies. A dedicated catalog volume pool facilitates quicker catalog restore times. Multiple catalog backup volume pools are allowed.
Description	a brief description of the volume pool.
Maximum number of partially full media	This property does not apply to the None pool, catalog backup pools, or scratch volume pools. Specifies the number of partially full media to allow in the volume pool for each of the unique combinations of the following in that pool: ■ Robot ■ Drive type ■ Retention level The default value is zero, which does not limit the number of full media that are allowed in the pool.
Prefer span to scratch	Specifies how NetBackup should select additional media when tape media operations span multiple media. When this parameter is set to <code>yes</code> (default) if a job spans to new media, NetBackup selects media from the scratch pool. NetBackup takes this action instead of using partially full media from the backup volume pool. When this parameter is set to <code>no</code>, NetBackup attempts to select partially full media from the backup volume pool to complete the specified operation. The <code>no</code> setting lets NetBackup use partially full media in the backup volume pool instead of always spanning to a scratch tape. Set the maximum number of partially full media option with the <code>vmppool -create</code> or the <code>vmppool -update</code> command.
Pool name	The Pool name is the name for the new volume pool. Volume pool names are case-sensitive and can be up to 20 characters.
Scratch pool	Specifies that the pool should be a scratch pool. It is recommended that you use a descriptive name for the pool and use the term <code>scratch pool</code> in the description. Add sufficient type and quantity of media to the scratch pool to service all scratch media requests that can occur. NetBackup requests scratch media when media in the existing volume pools are allocated for use. In NetBackup there can be only one scratch pool

Managing volume groups

You can perform the following tasks to manage volume groups.

See [“Delete a volume group”](#) on page 319.

See [“Move a volume group”](#) on page 319.

Delete a volume group

Use the following procedure to delete a volume group.

To delete a volume group

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Then select the **Volume groups** tab.
- 3 In the volumes list, verify that all of the volumes in the group are unassigned. You cannot delete the group until the application unassigns the volumes. If the **Time assigned** column contains a value, the volume is assigned.
- 4 Select one or more volume groups to delete.
- 5 Select **Delete**.
- 6 Select **Confirm**.
- 7 Remove the deleted volumes from the storage device.

Note: After you delete a volume group, it deletes the volumes in the volume groups.

Move a volume group

You can move a volume group from a robotic library to standalone storage or from standalone storage to a robotic library.

Moving a volume group changes only the residence information in NetBackup. You must move the volumes physically to their new locations.

To move a volume group

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage**. Then select the **Volume groups** tab.
- 3 Select the volume group that you want to move.
- 4 Select **Move**.

- 5 You can only specify the properties that apply for the move type.

Property	Description
Selected volume group	The volume group to move.
Robot	When you move from a robotic library, this value displays the robot type, robot number, and robot control host. When you move standalone volumes, this value displays "Standalone".
Destination	If you move a volume group to a robotic library, select the Device host and the Robot. If you move a volume group to standalone storage, no configuration is needed.
Device host	The host that controls the robotic library.
Robot	The destination robotic library.

- 6 Select **Confirm**.
- 7 After you move the volume group logically, physically move the volumes to their new locations.

Inventorying robots

This chapter includes the following topics:

- [About robot inventory](#)
- [When to inventory a robot](#)
- [About showing a robot's contents](#)
- [Show the media in a robot](#)
- [About comparing a robot's contents with the volume configuration](#)
- [Comparing media in a robot with the volume configuration](#)
- [About previewing volume configuration changes](#)
- [Previewing volume configuration changes for a robot](#)
- [About updating the NetBackup volume configuration](#)
- [Update the NetBackup volume configuration with a robot's contents](#)
- [Robot inventory options](#)
- [Advanced options for robot inventory settings](#)
- [Configure media ID generation rules](#)
- [Barcode rules settings](#)
- [Media ID generation options](#)
- [Configure media settings](#)
- [About media type mapping rules](#)
- [Configure media type mappings](#)

About robot inventory

Robot inventory is a logical operation that verifies the presence of media. (Robot inventory does not inventory the data on the media.)

After you physically add, remove, or move volumes in a robot, use a robot inventory to update the NetBackup volume configuration.

The following table describes the robot inventory options for the robotic libraries that contain barcode readers and contain barcoded media.

Table 16-1 Robot inventory options

Inventory option	Description
Show contents	Queries the robot for its contents and displays the media in the selected robotic library; does not check or change the EMM database. See “About showing a robot's contents” on page 324. For the robotic libraries without barcode readers (or that contain media without barcodes), you can only show the contents of a robot. However, more detailed information is required to perform automated media management. Use the <code>vmphyinv</code> physical inventory utility to inventory such robots.
Compare contents with volume configuration	Queries the robot for its contents and compares the contents with the contents of the EMM database. This option does not change the database. See “About comparing a robot's contents with the volume configuration” on page 326.
Preview volume configuration changes	Queries the robot for its contents and compares the contents with the contents of the EMM database. If differences exist, it is recommended to change to the NetBackup volume configuration. See “About previewing volume configuration changes” on page 327.
Update volume configuration	Queries the robot for its contents; if necessary, updates the database to match the contents of the robot. If the robot contents are the same as the EMM database, no changes occur. See “About updating the NetBackup volume configuration” on page 328.

When to inventory a robot

The following table describes the criteria to use to determine when to inventory a robot and which options to use for the inventory.

Table 16-2 Robot inventory criteria

Action	Inventory option to use
To determine the contents of a robot	Use the Show contents option to determine the media in a robot and possibly their barcode numbers. See “Show the media in a robot” on page 325.
To determine if volumes were moved physically within a robot	For the robots with barcode readers and the robots that contain media with barcodes, use the Compare contents with volume configuration option. See “Comparing media in a robot with the volume configuration” on page 326.
To add new volumes to a robot (a new volume is one that does not have a NetBackup media ID)	For any robot NetBackup supports, use the Update volume configuration option. The update creates media IDs (based on barcodes or a prefix that you specify). See “Update the NetBackup volume configuration with a robot's contents” on page 329.
To determine whether new media have barcodes before you add them to NetBackup	Use the Preview volume configuration changes option, which compares the contents of the robot with the NetBackup volume configuration information. After you examine the results, use the Update volume configuration option to update the volume configuration if necessary. See “Update the NetBackup volume configuration with a robot's contents” on page 329.
To insert existing volumes into a robot (an existing volume is one that already has a NetBackup media ID)	If the robot supports barcodes and the volumes have readable barcodes, use the Update volume configuration option. NetBackup updates the residence information to show the new robotic location. NetBackup also updates the robot host, robot type, robot number, and slot location. Specify the volume group to which the volume is assigned. See “Update the NetBackup volume configuration with a robot's contents” on page 329. If the robot does not support barcodes or the volumes do not contain readable barcodes, move the volumes or use the physical inventory utility.
To move existing volumes between robotic and standalone (an existing volume is one that already has a NetBackup media ID)	If the robotic library supports barcodes and the volumes have readable barcodes, use the Update volume configuration option. NetBackup updates the residence information to show the new robotic or standalone location. See “Update the NetBackup volume configuration with a robot's contents” on page 329.

Table 16-2 Robot inventory criteria (*continued*)

Action	Inventory option to use
To move existing volumes within a robot (an existing volume is one that already has a NetBackup media ID)	If the robot supports barcodes and the volumes have readable barcodes, use the Update volume configuration option. NetBackup updates the residence information to show the new slot location. See “Update the NetBackup volume configuration with a robot's contents” on page 329. If the robot does not support barcodes or if the volumes do not contain readable barcodes, move the volumes or use the physical inventory utility.
To move existing volumes from one robot to another (an existing volume is one that already has a NetBackup media ID)	If the robotic library supports barcodes and the volumes have readable barcodes, use the Update volume configuration option. NetBackup updates the NetBackup volume configuration information. See “Update the NetBackup volume configuration with a robot's contents” on page 329. If the robots do not support barcodes or the volumes do not contain readable barcodes, move the volumes or use the physical inventory utility. For either operation, perform the following updates: ■ First move the volumes to standalone ■ Then move the volumes to the new robot If you do not perform both updates, NetBackup cannot update the entries and writes an "Update failed" error.
To remove existing volumes from a robot (an existing volume is one that already has a NetBackup media ID)	For any robot NetBackup supports, use the Update volume configuration option to update the NetBackup volume configuration information. See “Update the NetBackup volume configuration with a robot's contents” on page 329.

About showing a robot's contents

Show contents inventories the selected robotic library and generates a report. This operation does not check or change the EMM database. Use this option to determine the contents of a robot.

The contents that appear depend on the robot type.

The following table describes the report contents.

Note: On UNIX: If a volume is mounted in a drive, the inventory report lists the slot from which the volume was moved to the drive.

Table 16-3 Show contents description

Robot and media	Report contents
The robot has a barcode reader and the robot contains media with barcodes.	Shows if each slot has media and lists the barcode for the media.
The robot does not have a barcode reader or the robot contains media without barcodes.	Shows if each slot has media.
API robot.	Shows a list of the volumes in the robot.

See [“Show the media in a robot”](#) on page 325.

About inventory results for API robots

The following table describes the contents of the robot inventory for the API robots.

Table 16-4 API robot report contents

Robot type	Report contents
ACS	The results, received from ACS library software, show the following: ■ The ACS library software volume ID. The NetBackup media ID corresponds to the ACS library software volume ID. ■ The ACS media type. ■ The NetBackup Media Manager media type. ■ The mapping between the ACS library software media type and the corresponding NetBackup Media Manager media type (without considering optional barcode rules).

Show the media in a robot

Use the following procedure to show the media that is in a robot.

See [“About robot inventory”](#) on page 322.

See [“Robot inventory options”](#) on page 329.

To show the media in a robot

- 1 Open the NetBackup web UI.
- 2 On the left, select **Storage > Tape storage > Robots**.
- 3 Select the robot that you want to inventory.
- 4 Select **Inventory robot**.

- 5 Ensure that the correct **Device host** and **Robot** are selected.
- 6 Go to **Inventory operation** and select **Show contents**.
- 7 Select **Start** to begin the inventory.

About comparing a robot's contents with the volume configuration

Compare contents with volume configuration compares the contents of a robotic library with the contents of the EMM database. Regardless of the result, the database is not changed.

Table 16-5 Compare contents description

Robot and media	Report contents
The robot can read barcodes	The report shows the differences between the robot and the EMM database
The robot cannot read barcodes	The report shows only whether a slot contains a volume If the media have barcodes, this operation is useful for determining if volumes have been physically moved within a robot.
For API robots	The media ID and media type in the EMM database are compared to the information that is received from the vendor's robotic library software.

If the results show that the EMM database does not match the contents of the robotic library, perform the following actions:

- Physically move the volume.
- Update the EMM database.

See [“About updating the NetBackup volume configuration”](#) on page 328.

See [“Comparing media in a robot with the volume configuration”](#) on page 326.

Comparing media in a robot with the volume configuration

Use the following procedure to compare the media in a robot with the EMM database.

See [“About robot inventory”](#) on page 322.

See [“Robot inventory options”](#) on page 329.

To compare media in a robot with the volume configuration

- 1 In the **NetBackup web UI**, click **Storage > Tape storage > Robots**.
- 2 Select the robot that you want to inventory.
- 3 On the **Actions** menu, select **Inventory Robot**.
- 4 In the **Inventory options**, select **Compare contents with volume configuration**.
- 5 Click **Start** to begin the inventory.

About previewing volume configuration changes

Use this option to preview the changes before you update the EMM database. This option lets ensure that all new media have barcodes before you add them to the EMM database.

Note: If you preview the configuration changes first and then update the EMM database, the update results may not match the results of the preview operation. Possible causes may be the changes that occur between the preview and the update. Changes can be to the state of the robot, to the EMM database, to the barcode rules, and so on.

See [“Update the NetBackup volume configuration with a robot’s contents”](#) on page 329.

Previewing volume configuration changes for a robot

Use the procedure in this topic to preview any volume configuration changes for a robot.

See [“About previewing volume configuration changes”](#) on page 327.

See [“Robot inventory options”](#) on page 329.

To preview the volume configuration changes for a robot

- 1 If necessary, add new volumes into the robotic library.
- 2 In the **NetBackup web UI**, click **Storage > Tape storage > Robots**.

- 3 Select the robot you want to inventory.
- 4 On the **Actions** menu, select **Inventory Robot**.
- 5 On the **Inventory operations**, select **Preview volume configuration changes**.

Note: If you preview the configuration changes first and then update the EMM database, the update results may not match the results of the preview operation. Possible causes may be the changes that occur between the preview and the update. Changes can be to the state of the robot, to the EMM database, to the barcode rules, and so on.

- 6 To change the default settings and rules that NetBackup uses to name and assign attributes to new media, click **Advanced options**.

Note: Advanced options apply to only preview and update volume configuration and hence enabled only when you select these operation options.

- 7 To inject any media that is in the media access port before the preview operation, click **Empty media access port prior to update**.
- 8 Click **Start** to begin the inventory preview.

About updating the NetBackup volume configuration

The **Update volume configuration** robot inventory option updates the database to match the contents of the robot. If the robot contents are the same as the EMM database, no changes occur.

For a new volume (one that does not have a NetBackup media ID), the update creates a media ID. The media ID depends on the rules that are specified on the **Advanced options** section.

See [“Robot inventory options”](#) on page 329.

For API robots, the update returns an error if the volume serial number or the media ID contain unsupported characters.

For robots without barcode readers, the new media IDs are based on a media ID prefix that you specify. Similarly, for volumes without readable barcodes, the new media IDs are based on a media ID prefix that you specify.

Robot inventory update returns an error if it encounters unsupported characters in the volume serial number or media identifier from API robots.

See [“Update the NetBackup volume configuration with a robot's contents”](#) on page 329.

Update the NetBackup volume configuration with a robot's contents

Use the procedure in this topic to update the EMM database with the contents of a robot.

See [“About updating the NetBackup volume configuration”](#) on page 328.

See [“Robot inventory options”](#) on page 329.

To update the volume configuration with a robot's contents

- 1 If necessary, add new volumes into the robotic library.
- 2 Open the web UI.
- 3 On the left, select **Storage > Tape storage**. Select the **Robots** tab.
- 4 Select the robot that you want to inventory.
- 5 Select **Actions > Inventory robot**.
- 6 Select **Update volume configuration**.

Note: If you preview the configuration changes first and then update the EMM database, the update results may not match the results of the preview operation. Possible causes may be the changes that occur between the preview and the update. Changes can be to the state of the robot, to the EMM database, to the barcode rules, and so on.

- 7 To change the default settings and rules that NetBackup uses to name and assign attributes to new media, select **Advanced options**.
- 8 To inject any media that is in the media access port before the update operation, click **Empty media access port prior to update**.
- 9 Select the **Start** button to begin the inventory update.

Robot inventory options

The following table shows the robot inventory options.

Table 16-6 Robot inventory options

Option	Description
Advanced options	Advanced options is active if Preview volume configuration changes or Update volume configuration is selected. This button opens the Advanced robot inventory options tab from which you can configure more options. See “Advanced options for robot inventory settings” on page 331.
Device host	The Device host option is the host that controls the robot.
Empty media access port prior to update	The Empty media access port prior to update operation is active only for the robots that support that function. To inject volumes in the robot’s media access port into the robot before you begin the update, select Empty media access port prior to update. The volumes to be injected must be in the media access port before the operation begins. If you select Empty media access port prior to update and the media access port is empty, you are not prompted to place volumes in the media access port. Note: If you use NetBackup to eject volumes from the robot, remove the volumes from the media access port before you begin an inject operation. Otherwise, if the inject port and eject port are the same, the ejected volumes may be injected back into the robotic library.
Robot	Use the Robot option to select a robot to inventory. If you selected a robot in the NetBackup web UI, that robot appears in this field.
Show contents	Displays the media in the selected robotic library; does not check or change the EMM database. See “About showing a robot’s contents” on page 324.
Compare contents with volume configuration	Compares the contents of a robotic library with the contents of the EMM database but does not change the database. See “About comparing a robot’s contents with the volume configuration” on page 326.
Preview volume configuration changes	Compares the contents of a robotic library with the contents of the EMM database. If differences exist, it is recommended to change to the NetBackup volume configuration. See “About previewing volume configuration changes” on page 327.

Table 16-6 Robot inventory options (*continued*)

Option	Description
Update volume configuration	Updates the database to match the contents of the robot. If the robot contents are the same as the EMM database, no changes occur. See “About updating the NetBackup volume configuration” on page 328.

Table 16-7 Results pane

Option	Description
Show contents	Displays the contents of robot (Slot, Tape, Barcode). Note: While you are running the robot inventory contents results job, if the result is taking time, you can leave the page and return to find while the complete result is displayed. See “About showing a robot's contents” on page 324.
Compare contents	Displays the comparison between the Robot Contents (Slot, Tape, and Barcode) and Volume Configuration (Media ID, and Barcode) with the Mismatch Detected list. See “About comparing a robot's contents with the volume configuration” on page 326.
Preview volume configuration changes	Lists the proposed changes to EMM database Volume Configuration). To update the volume configuration changes. See “About previewing volume configuration changes” on page 327.
Update	Lists the updated changes as well as the actual changes performed, along with the success message. See “About updating the NetBackup volume configuration” on page 328.
Download	In the case, the robot inventory result text is large (more than 100K results) the web UI shows the truncated data with an option to download the text file.
Search	Allows a search for specific term/keyword in the results text.
Copy to clipboard	Allows to copy the results text to the clipboard.

Advanced options for robot inventory settings

The following advanced options are available for robot inventory.

Table 16-8 Advanced options for robot inventory settings

Option	Description
Media settings	
Existing media	Media which have been removed from the robot should be assigned to the volume group options: ■ Default: If there is an existing group with a compatible residence for the volume, the volume is added to that group. If a suitable volume group does not exist, NetBackup generates a new volume group name. ■ Auto-generate: NetBackup automatically generates a new volume group. ■ No volume group: The media are not assigned to a volume group. Media which have been moved into or within the robot should be assigned to the volume group options: ■ Default: Includes the volume groups that are valid for the robot's default media type. ■ Auto-generate: NetBackup automatically generates a new volume group. ■ If the Media type is a value other than Default: Includes the volume groups that are valid for the specified media type. To specify a volume group other than Default, select a volume group name from the list.

Table 16-8 Advanced options for robot inventory settings (*continued*)

Option	Description
New media	Use barcode rules Specifies whether or not to use barcode rules to assign attributes for new media. To enable barcode rule support for API robots, add an <code>API_BARCODE_RULES</code> entry to the <code>vm.conf</code> file. Media type Overrides your barcode rules for the new media in the robotic library. Volume pool Overrides the default volume pool for the new media in the robotic library. Determine how to use a Media ID prefix: ■ To not use a media ID prefix: Deselect the Use a media ID prefix for media with unreadable barcodes or if the robot does not support barcodes option. ■ To use a media ID prefix: Select the Use a media ID prefix for media with unreadable barcodes or if the robot does not support barcodes option. ■ To use a specific media ID prefix for the current session only: Select the Specify the media ID prefix for the current session only option then enter the media ID prefix. You can specify a prefix of one to five alphanumeric characters. NetBackup assigns the remaining numeric characters to create a six character media ID. NetBackup uses the prefix only for the current operation. ■ Media ID prefix support only A-Z, 0-9, ‘_’ characters. Underscore ‘_’ not allowed as first character. ■ Choose from the media ID prefix list (stored in <code>vm.conf</code> file): Select this option and then enter the prefix. Click Add to the list.
Barcode rules	
Add	Click Add to add a new barcode rule. See “Barcode rules settings” on page 335.
Media ID generation	

Table 16-8 Advanced options for robot inventory settings (*continued*)

Option	Description
Add Media ID generation	Media ID generation rules let you override the default media ID naming method. The default method generates a media ID using the last six characters of the barcode. Click Add to add a new rule. See " Media ID generation options " on page 337.

Configure media ID generation rules

For non-API robots only. Robot types are described in a different topic.

Use the **Media ID generation** option to configure the rules that override the default naming method. To use media ID generation rules, the robot must support barcodes and the robot cannot be an API robot.

To configure media ID generation rules

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**.
- 3 On the top right, click **Actions > Inventory robots**.
- 4 Select the **Device host**.
- 5 Select the robot.
- 6 Click either **Preview volume configuration changes** or **Update volume configuration**.
- 7 Click on **Advanced options**. Then click **Media ID generation**.

8 To configure the rules, do any of the following:

- Add a rule

Click **Add** and then configure the rule.
- Edit a rule

Locate the rule and click **Edit**.

You cannot change the robot number or barcode length of a rule.
To change those properties, first delete the old rule and then add a rule.
- Delete a rule

Locate the rule and click **Delete**.

9 When you are finished configuring rules, click **Save**.

Note: If you click save on the individual rows does NOT save the rule, only when the you clicks the dialog's save button are all the changes saved.

Barcode rules settings

The following table describes the settings you can configure for barcode rules. NetBackup uses these rules to assign barcodes to new media.

Table 16-9 Barcode rule settings

Barcode rule setting	Description
Barcode tag	A unique string of barcode characters that identifies the type of media. For example, use DLT as the barcode tag for a barcode rule if the following is true: ■ You use DLT on the barcodes to identify DLT tapes ■ DLT is not used on any other barcodes in the robot Similarly, if you use CLND for DLT cleaning media, use CLND as the barcode tag for the rule for DLT cleaning media. The barcode tag can have from 1 to 16 characters but cannot contain spaces. The following are the special barcode rules that can match special characters in the barcode tags: ■ NONE Matches when rules are used and volume has an unreadable barcode or the robot does not support barcodes. ■ Barcode rule names support only alphabets A-Z, numerics 0-9, special character Underscore '_' for barcode rule names. Underscore '_' not allowed as first character. You can change/edit a barcode tag of a barcode rule in the web UI. Use the Media Settings tab to set up the criteria for a robot update.
Description	A description of the barcode rule. Enter from 1 to 25 characters.
Maximum mounts	The maximum number of mounts (or cleanings) that are allowed for the volume. For data volumes, a value of zero means the volume can be mounted an unlimited number of times. For cleaning tapes, zero means that the cleaning tape is not used. It is recommended that you use barcodes for the cleaning media that cannot be confused with barcodes for data media. The media type to assign to the media. Doing so can avoid a value of 0 for cleaning tapes.
Media type option	If media type selected is a cleaning tape then Volume pool is not selectable and set to None.

Table 16-9 Barcode rule settings (*continued*)

Barcode rule setting	Description
Volume pool	The volume pool for the new media. The actions depend on whether you use barcode rules to assign media attributes. Select from the following: ■ DEFAULT If DEFAULT is selected, NetBackup performs the following actions: ■ If you use barcode rules, the barcode rules determine the volume pool to which new volumes are assigned. ■ If you do not use barcode rules, NetBackup assigns data tapes to the NetBackup pool but does not assign cleaning tapes to a volume pool. ■ A specific volume pool This volume pool setting always overrides any barcode rules.

Media ID generation options

NetBackup uses rules to generate the IDs for media in robots. The default rule uses the last six characters of the barcode label from the tape.

You can configure media ID generation rules to override the default rule. Control how NetBackup creates media IDs by defining the rules that specify which characters of a barcode label to use for the media ID.

The following subsections describe the media ID generation rule options.

The following list describes the media ID generation rule options:

- **Bar code length**
The **Barcode length** is the number of characters in the barcode for tapes in the robot.
You cannot change the barcode length of a rule. Rather, first delete the rule and then add a new rule.
- **Media ID generation rule**
A **Media ID generation rule** consists of a maximum of six colon-separated fields. Numbers define the positions of the characters in the barcode that are to be extracted. For example, the number 2 in a field extracts the second character (from the left) of the barcode. You can specify numbers in any order.
To insert a specific character in a generated media idea, precede the character by a pound sign (#). Any alphanumeric characters that are specified must be valid for a media ID.
Use rules to create media IDs of many formats. Ensure that the media ID generation rule generates a unique media ID.

A media ID generation rule is a rule to convert a barcode (which can be up to 16 characters) into a media id (which is limited to 6 characters).

The rule allows you to specify literal characters, or character positions from the original barcode.

Characters of the rule are separated by ':'. A number represents a character in the barcode, starting at 1.

A '#' followed by a character represents a literal character.

For example, #A:3:2:1 would result in the letter A followed by the third, second, and then first characters of * the volume's barcode.

No two rules can share both robot number and barcode length.

The table shows some examples of rules and the resulting media IDs.

Barcode on tape	Media ID generation rule	Generated media ID
032945L1	1:2:3:4:5:6	032945
032945L1	3:4:5:6:7	2945L
032945L1	#N:2:3:4:5:6	N32945
543106L1	#9:2:3:4	9431
543106L1	1:2:3:4:#P	5431P

■ Robot number

The number of the robot to which the rule applies.

You cannot change the robot number of a rule. Rather, first delete the rule and then add a new rule.

Configure media settings

This procedure describes how to configure the attributes for existing and new media.

To configure media settings

- 1 Open the NetBackup web UI.
- 2 On the left, click **Storage > Tape storage**.
- 3 On the top right, click **Actions > Inventory robots**.
- 4 Select the **Device host**.
- 5 Select the robot.
- 6 Click either **Preview volume configuration changes** or **Update volume configuration**.

- 7 Click on **Advanced options**. Then click **Media settings**.
- 8 Configure the settings as follows:
 - a. In the **Media which have been removed from the robot should be assigned to the volume group** list, select a volume group for the media that are removed from the robot.

See [“Advanced options for robot inventory settings”](#) on page 331.
 - b. In the **Media which have been moved into or within the robot should be assigned to the volume group** list, select a volume group for the media that are in or are added to the robot.

See [“Advanced options for robot inventory settings”](#) on page 331.
 - c. If the robotic library supports barcodes and the volume has readable barcodes, NetBackup creates media IDs automatically from the barcodes. You do not need to configure a prefix.

However, if the media in the robotic library has unreadable barcodes or if the robot does not support barcodes, NetBackup assigns a default media ID prefix.

To use a media ID prefix other than the **Default**, click **Use the following Media ID prefix** field. Then, specify or choose a media ID prefix.

See [“Advanced options for robot inventory settings”](#) on page 331.
 - d. To use your barcode rules to assign attributes to new volumes, select **Use barcode rules**.

See [“Advanced options for robot inventory settings”](#) on page 331.
 - e. To override your barcode rules for the new media in the robotic library, select a **Media type** from the list.

See [“Advanced options for robot inventory settings”](#) on page 331.
 - f. To override the default volume pool for the new media in the robotic library, select a **Volume pool** from the list.

See [“Advanced options for robot inventory settings”](#) on page 331.
- 9 Click **Save**.

About media type mapping rules

Applies to API robots only. Robot types are described in a different topic.

For API robots, NetBackup contains default mappings from a vendor's media types to NetBackup media types. API robots are ACS robot types.

You can change the default mappings. Changes apply only to the current volume configuration update.

You can also add media type mappings.

Note: You can write a barcode rule that contains the media types that are incompatible with vendor media types. However, the robot inventory update may assign NetBackup media types that are inconsistent with the vendor media types. Avoid this problem by grouping barcode rules by media type.

Configure media type mappings

Use the **Media type mappings** in the Advanced options for robot inventory to configure the attributes for existing and new media.

See [“About media type mapping rules”](#) on page 339.

To configure media type mappings

- 1 Open the web UI.
- 2 On the left, select **Storage > Tape storage**.
- 3 At the top right, select **Actions > Inventory robot**.
- 4 Select the **Device host**. Then select the **Robot** that you want to inventory.
- 5 Select either **Preview volume configuration changes** or **Update volume configuration**.
- 6 Select **Advanced options** and select **Media type mappings**.

The **Media Type mappings** are only available for if the **Robot type** is **ACS**.

The mappings that appear are only for the robot type that was selected for inventory. The default mappings and any mappings that you added or changed appear.

- 7 Locate the row that contains the robot-vendor media type mapping that you want to change and select **Edit**.
- 8 Select a **Media type** from the list.
- 9 Select **Save**.

Staging backups

This chapter includes the following topics:

- [About staging backups](#)
- [About basic disk staging](#)
- [Create a BasicDisk storage unit with disk staging](#)
- [Disk staging storage unit size and capacity](#)
- [Finding potential free space on a BasicDisk disk staging storage unit](#)
- [Schedule settings for disk staging](#)

About staging backups

In the staged backups process, NetBackup writes a backup to a storage unit and then duplicates it to a second storage unit. Eligible backups are deleted on the initial storage unit when space is needed for more backups.

This two-stage process allows a NetBackup environment to leverage the advantages of disk-based backups for recovery in the short term.

Staging also meets the following objectives:

- Allows for faster restores from disk.
- Allows the backups to run when tape drives are scarce.
- Allows the data to be streamed to tape without image multiplexing.

NetBackup offers the following methods for staging backups.

Table 17-1 Methods for staging backups

Staging method	Description
Basic disk staging	Basic disk staging consists of two stages. First, data is stored on the initial storage unit (disk staging storage unit). Then, per a configurable relocation schedule, data is copied to the final location. Having the images on the final destination storage unit frees the space on the disk staging storage unit as needed. See “About basic disk staging” on page 342. The following storage unit types are available for basic disk staging: BasicDisk and tape.
Staging using the Storage lifecycle policies utility	Staged backups that are configured within the Storage lifecycle policies utility also consist of two stages. Data on the staging storage unit is copied to a final destination. However, the data is not copied per a specific schedule. Instead, the administrator can configure the data to remain on the storage unit until either a fixed retention period is met, or until the disk needs additional space, or until the data is duplicated to the final location. No BasicDisk or disk staging storage unit can be used in an SLP.

About basic disk staging

Basic disk staging is conducted in the following stages.

Table 17-2 Basic disk staging

Stage	Description
Stage I	Clients are backed up by a policy. The Policy storage selection in the policy indicates a storage unit that has a relocation schedule configured. The schedule is configured in the staging schedule settings.
Stage II	Images are copied from the Stage I disk staging storage unit to the Stage II storage unit. The relocation schedule on the disk staging storage unit determines when the images are copied to the final destination. Having the images on the final destination storage unit frees the space on the disk staging storage unit as needed.

The image continues to exist on both the disk staging storage unit and the final destination storage units until the image expires or until space is needed on the disk staging storage unit.

When the relocation schedule runs, NetBackup creates a data management job. The job looks for any data that can be copied from the disk staging storage unit to the final destination. The job details in the Activity monitor identify the job as one associated with basic disk staging. The jobs list displays Disk Staging in the job’s **Data movement** field.

When NetBackup detects a disk staging storage unit that is full, it pauses the backup. Then, NetBackup finds the oldest images on the storage unit that successfully copied onto the final destination. NetBackup expires the images on the disk staging storage unit to create space.

Note: The basic disk staging method does not support backup images that span disk storage units.

To avoid spanning storage units, do not use Checkpoint restart on a backup policy that writes to a storage unit group that contains multiple disk staging storage units.

Create a BasicDisk storage unit with disk staging

When you configure a BasicDisk storage unit with disk staging, data is stored on the initial storage unit (disk staging storage unit). Then, per a configurable relocation schedule, data is copied to the final location. Having the images on the final destination storage unit frees the space on the disk staging storage unit as needed.

To create a BasicDisk storage unit with disk staging

- 1 Click **Storage > Storage units**.
- 2 Click **Add**.
- 3 Select **BasicDisk**. Then click **Start**.
- 4 Select the basic properties for the storage unit.

Type a **Name** for the storage unit.

Enter the number of **Maximum concurrent jobs** that are allowed to write to this storage unit at one time.

Enter a **High water mark** value.

The high water mark works differently for the BasicDisk disk type. NetBackup assigns new jobs to a BasicDisk disk staging storage unit, even if it is over the indicated high water mark. For BasicDisk, the high water mark is used to prompt the deletion of images that have been relocated.

Note: The **Low water mark** setting does not apply to disk staging storage units.

- 5 Click **Next**.

- 6 For the staging schedule, select the option **Enable temporary staging area**.
- 7 Below **Staging schedule**, click **Add**.

The schedule name defaults to the storage unit name.

Configure the schedule settings.

See [“Schedule settings for disk staging”](#) on page 347.
- 8 Click **Save** to save the disk staging schedule.
- 9 Click **Next**.
- 10 Select a media server.
- 11 Browse or specify the absolute path to the directory to be used for storage.
- 12 Select whether this directory can reside on the root file system or system disk.
- 13 Click **Next**.
- 14 Review the settings for the storage unit and then click **Save**.

Disk staging storage unit size and capacity

To take advantage of basic disk staging requires that the NetBackup administrator understand the life expectancy of the image on the Stage I storage unit.

The size and use of the file system of the Stage I storage unit directly affects the life expectancy of the image before it is copied to the Stage II storage unit. It is recommended a dedicated file system for each disk staging storage unit.

Consider the following example: A NetBackup administrator wants incremental backups to be available on disk for one week.

Incremental backups are done Monday through Saturday, with full backups done on Sunday. The full backups are sent directly to tape and do not use basic disk staging.

Each night's total incremental backups are sent to a disk staging storage unit and average from 300 MB to 500 MB. Occasionally a backup is 700 MB. Each following day the relocation schedule runs on the disk staging storage unit and copies the previous night's incremental backups to the final destination, a Media Manager (tape) storage unit.

The following items give more information about determining disk size for a basic disk staging storage unit.

Minimum disk size

The minimum disk size is the smallest size that is required for the successful operation of the disk staging logic.

The minimum size must be greater than or equal to the largest combined size of the backups that are placed on the storage unit between runs of the disk staging schedule. (In our example, the disk images remain on the disk for one week.)

In this example, the relocation schedule runs nightly, and the largest nightly backup is 700 MB. It is recommended that you double this value to allow for any problems that may occur when the relocation schedule runs. To double the value gives the administrator an extra schedule cycle (one day) to correct any problems.

To determine the minimum size for the storage unit in this example, use the following formula:

Minimum size = Max data per cycle × (1 cycle + 1 cycle for safety)

For example: 1.4 GB = 700 MB × (1+1)

Average disk size

The average disk size represents a good compromise between the minimum and the maximum sizes.

In this example, the average nightly backup is 400 MB and the NetBackup administrator wants to keep the images for one week.

To determine the average size for the storage unit in this example, use the following formula:

Average size = Average data per cycle × (number of cycles to keep data + 1 cycle for safety)

2.8 GB = 400 MB × (6 + 1)

Maximum disk size

The maximum disk size is the recommended size needed to accommodate a certain level of service. In this example, the level of service is that disk images remain on disk for one week.

To determine the maximum size for the storage unit in this example, use the following formula:

Maximum size = Max data per cycle × (# of cycles to keep data + 1 cycle for safety)

For example: 4.9 GB = 700 MB × (6 + 1)

Finding potential free space on a BasicDisk disk staging storage unit

Potential free space is the amount of space on a disk staging storage unit that NetBackup could free if extra space on the volume is needed. The space is the

total size of the images that are eligible for expiration plus the images ready to be deleted on the volume.

To find the potential free space on a BasicDisk storage unit, use the `bpstulist` and the `nbdevquery` commands as follows:

- Run `bpstulist -label` to find the disk pool name.
Note that the name of the storage unit and disk pools are case-sensitive. In the case of BasicDisk storage units, the name of the disk pool is the same as the name of the BasicDisk storage unit. In the following example, the name of the storage unit is *NameBasic*:

```
bpstulist -label basic
NameBasic 0 server1 0 -1 -1 1 0 "C:\\" 1 1 524288 *NULL* 0 1 0 98 80 0 NameBasic server1
```

- Run the `nbdevquery` command to display the status for the disk pool, including the potential free space.
Use the following options, where:

<code>-stype server_type</code>	Specifies the vendor-specific string that identifies the storage server type. For a BasicDisk storage unit, enter <i>BasicDisk</i> .
<code>-dp</code>	Specifies the disk pool name. For a basic disk type, the disk pool name is the name of the BasicDisk storage unit.

So the complete command might look like the following.

```
nbdevquery -listdv -stype BasicDisk -dp NameBasic -D
```

The value is listed as `potential_free_space`.

```
Disk Volume Dump
name           : <Internal_16>
id             : <C:\>
diskpool       : <NameBasic::server1::BasicDisk>
disk_media_id  : <@aaaaaf>
total_capacity : 0
free_space     : 0
potential_free_space: 0
committed_space : 0
precommitted_space : 0
nbu_state      : 2
sts_state      : 0
```

```
flags                : 0x6
num_read_mounts      : 0
max_read_mounts      : 0
num_write_mounts     : 1
max_write_mounts     : 1
system_tag           : <Generic disk volume>
```

Schedule settings for disk staging

The following settings are available when you create a disk staging schedule.

Table 17-3 The Attributes tab settings

Attribute	Description
Name	The schedule Name defaults to the name of the storage unit.
Priority of relocation jobs started from this schedule	The Priority of relocation jobs started from this schedule field indicates the priority that NetBackup assigns to relocation jobs for this policy. Range: 0 to 99999 (highest priority). The default value that is displayed is the value that is set in the Default job priorities host properties for the Staging job type.
Multiple copies	Creates multiple copies of backups. NetBackup can create up to four copies of a backup simultaneously. When this setting is enabled, Final destination volume pool and Final destination media ownership are disabled.
Final destination storage unit	If the schedule is a relocation schedule, a Final destination storage unit must be indicated. (A relocation schedule is created as part of a basic disk staging storage unit configuration.) A Final destination storage unit is the name of the storage unit where the images reside after a relocation job copies them. To copy images to tape, NetBackup uses all of the drives available in the Final destination storage unit. However, the Maximum concurrent write drives setting for that storage unit must be set to reflect the number of drives. The setting determines how many duplication jobs can be launched to handle the relocation job. NetBackup continues to free space until the Low water mark is reached. See “About staging backups” on page 341.
Final destination volume pool	If the schedule is a relocation schedule, a Final destination volume pool must be indicated. (A relocation schedule is created as part of a basic disk staging storage unit configuration.) A Final destination volume pool is the volume pool where images are swept from the volume pool on the basic disk staging storage unit. See “About staging backups” on page 341.

Table 17-3 The Attributes tab settings (*continued*)

Attribute	Description
Final destination media owner	If the schedule is a relocation schedule, a Final destination media owner must be indicated. (A relocation schedule is created as part of a basic disk staging storage unit configuration.) A Final destination media owner is the media owner where the images reside after a relocation job copies them. Specify one of the following: ■ Any lets NetBackup choose the media owner. NetBackup chooses a media server or a server group (if one is configured). ■ None specifies that the media server that writes the image to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.
Schedule type	Calendar Frequency If the backups that use a disk staging storage unit run more frequently than expected, compare the retention level 1 setting with the Frequency setting. Internally, NetBackup uses the retention level 1 setting for scheduling purposes with disk staging storage units. Make sure that the frequency period is set to make the backups occur more frequently than the retention level 1 setting indicates. (The default is two weeks.) For example, a frequency of one day and a retention level 1 of 2 weeks should work well. Retention levels are configured in the Retention periods host properties.

Table 17-3 The Attributes tab settings (*continued*)

Attribute	Description
Use alternate read server	An alternate read server is a server allowed to read a backup image originally written by a different media server. The path to the disk or directory must be identical for each media server that is to access the disk. If the backup image is on tape, the media servers must share the same tape library or the operator must find the media. If the backup image is on a robot that is not shared or a standalone drive, the media must be moved to the new location. An administrator must move the media, inventory the media in the new robot, and run <code>bpmedia -oldserver -newserver</code> or assign a failover media server. To avoid sending data over the network during duplication, specify an alternate read server that meets the following conditions: ■ Connected to the storage device that contains the original backups (the source volumes). ■ Connected to the storage device that contains the final destination storage units. If the final destination storage unit is not connected to the alternate read server, data is sent over the network.
Copies	Specify the number of copies to create simultaneously. Range: 1 to 4.
Priority of duplication job	Indicates the priority that NetBackup assigns to duplication jobs for this policy. Range: 0 to 99999 (highest priority).

Table 17-3 The Attributes tab settings (*continued*)

Attribute	Description
Copy #	For each copy you want to create, select the copy settings. Copy 1 is the primary copy. If Copy 1 fails, the first successful copy is the primary copy. Storage unit Specify the storage unit where each copy is stored. If a Media Manager storage unit has multiple drives, it can be used for both the source and the destination. Volume pool Specify the volume pool where each copy is stored. If this copy fails ■ Continue Continues making the remaining copies. Note: Note: If Take checkpoints every __ minutes is selected for this policy, only the last failed copy that contains a checkpoint can be resumed. ■ Fail all copies Fails the entire job. Media owner For tape media, specify who should own the media onto which NetBackup writes the images. These settings do not affect any images that reside on disk. One media server does not own the images that reside on shared disks. Any media server with access to the shared pool of disk can access the images. ■ Any NetBackup selects the media owner, either a media server or server group. ■ None Specifies that the media server that writes to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.

Troubleshooting storage configuration

This chapter includes the following topics:

- [Registering a media server](#)
- [Storage configuration issues](#)

Registering a media server

If the primary server is not running when you install a media server, the media server is not registered. You cannot discover, configure, and manage the devices of that media server. You must register the media server with the primary server.

To register a media server

- 1 Start the EMM service on the primary server.
- 2 On the primary server, run the following command. (For *hostname*, use the host name of the media server.)

On Windows:

```
install_path\NetBackup\bin\admincmd\nbemcmd -addhost -machinename
hostname -machinetype media -masterserver server_name
-operatingsystem os_type -netbackupversion
level.major_level.minor_level
```

On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/nbemcmd -addhost -machinename
hostname -machinetype media -masterserver server_name
-operatingsystem os_type -netbackupversion
level.major_level.minor_level
```

Note: Ensure that the name you use in NetBackup is the same as the host name in the TCP/IP configuration.

Storage configuration issues

The following table describes multiple issues that might occur when you configure storage:

Table 18-1 Storage configuration troubleshooting

Error message or cause	Explanation and recommended action
The following error is displayed when you create a disk pool for a cloud volume: Disk is full	Workaround: Even if the disk is not full and you get the error, ensure that there is enough space available for creating the cloud volume. By default the cloud volume requires approximately 1 TB of free space. To reduce the cloud volume size, open the <code>contentrouter.cfg</code> file from <code>/msdp/etc/puredisk/</code> and change the values. After changing the values, restart the MSDP services and then create the cloud volume.
The local MSDP storage does not display the compression and the encryption values correctly.	In the Select long-term retention storage configuration page for protection plans, the local MSDP storage does not display the compression and the encryption values correctly.

Configuring backups

- [Chapter 19. Overview of backups in the NetBackup web UI](#)
- [Chapter 20. Managing classic policies](#)
- [Chapter 21. Managing protection plans](#)
- [Chapter 22. Protecting the NetBackup catalog](#)
- [Chapter 23. Managing backup images](#)
- [Chapter 24. Pausing data protection activity](#)

Overview of backups in the NetBackup web UI

This chapter includes the following topics:

- [Backup methods supported in the NetBackup web UI](#)
- [Policy vs. protection plan FAQs](#)
- [Support for NetBackup classic policies](#)
- [Supported protection plan types](#)

Backup methods supported in the NetBackup web UI

The NetBackup web UI offers the following methods to protect your data:

- **Policies.** Policies protect data on clients. Some agents also have an intelligent policy that protects assets spread over multiple clients.
- **Protection plans.** Protection plans protect assets. For example, databases or virtual machines. Workload administrators are granted access to a protection plans through the available default RBAC roles. Then they can subscribe the assets to a plan.

Protection plans and intelligent policies work with asset management to automatically discover assets in the NetBackup environment.

Policy vs. protection plan FAQs

You can protect an asset using a NetBackup classic policy, a protection plan, or both at the same time. This topic answers some common questions about NetBackup classic policies in the NetBackup web UI.

Table 19-1 Classic policy FAQ

Question	Answer
In the web UI's Protected by column, what does Classic policy only mean?	The asset is not currently subscribed to a protection plan. However, it was subscribed to a protection plan. Or, it was covered by a classic policy at one time and it has a Last backup status. There may or may not be an active classic policy protecting the asset (contact the NetBackup administrator to find out).
Where can I find the details of a classic policy?	The details of a classic policy are not visible in the web UI, with the exception of a few policy types. See “Support for NetBackup classic policies” on page 355.
How can I manage a classic policy?	Some policy types can be managed in the NetBackup web UI. See “Support for NetBackup classic policies” on page 355.
When should I subscribe an asset to a protection plan versus protecting the asset with a classic policy?	A protection plan lets you easily add and remove assets from the plan and see which assets are protected. A workload administrator can fully control who can view or manage protection plans and assets. Policies offer the classic method of data protection. However, they do not have RBAC control for individual policies or for the data you want to protect.
Can I use both a protection plan and a classic policy to protect an asset?	Yes. The web UI shows the details of the protection plan but not the details of the classic policy. You can contact the NetBackup administrator for the classic policy details.
What action should I take when an asset is unsubscribed from a protection plan and the web UI shows Classic policy only for that asset?	You can ask the NetBackup administrator if a classic policy protects the asset.

Support for NetBackup classic policies

The following policy types can be managed in the NetBackup web UI.

Table 19-2 Policy types supported in NetBackup web UI

BigData

Informix-On-BAR

NDMP

Table 19-2 Policy types supported in NetBackup web UI (*continued*)

Cloud (IaaS and PaaS)	Kubernetes	Nutanix-AHV
Cloud-Object-Store	Lotus Notes	Oracle
DataStore	MS-Exchange-Server	Teradata
DB2	MS-SharePoint	SAP
Enterprise Vault	MS-SQL-Server	Standard
Epic-Large-File	MS-Windows	Sybase
FlashBackup	MSDP-Object-Store	Universal-Share
FlashBackup-Windows	NAS-Data-Protection	VMware
Hyper-V	NBU-Catalog	

Supported protection plan types

The web UI supports protection plans for the following workloads.

Table 19-3

Apache Cassandra	OpenStack
Cloud	Oracle
Cloud object store	PostgreSQL
Kubernetes	Red Hat Virtualization (RHV)
Microsoft SQL Server	SaaS
MySQL	VMware
Nutanix AHV	

Managing classic policies

This chapter includes the following topics:

- [Add a policy](#)
- [Example policy - Exchange Server DAG backup](#)
- [Example policy - Sharded MongoDB cluster](#)
- [Example policy - Epic-Large-File](#)
- [Edit, copy, or delete a policy](#)
- [Deactivate or activate a policy](#)
- [View automanaged policies and SLPs](#)
- [About automanaged policies or storage lifecycle policies](#)
- [Perform manual backups](#)
- [About the Epic-Large-File policy type](#)

Add a policy

Use the following procedure to create a backup policy in the NetBackup web UI. Example policies are also available.

See [“Example policy - Exchange Server DAG backup”](#) on page 358.

See [“Example policy - Sharded MongoDB cluster”](#) on page 359.

For details on policy options, refer to the *NetBackup Administrator's Guide, Volume I* and to the appropriate workload or database guides.

Note: You must have the RBAC Administrator role or similar permissions to create and manage policies.

To add a policy

- 1 On the left, select **Protection > Policies**.
- 2 Click **Add**.
- 3 On the **Attributes** tab, do the following:
 - Specify the **Policy name**.
 - Select the **Policy type** that you want to create.
 - Select the **Policy storage** that you want to use.
 - Select or configure any other policy attributes.
- 4 On the **Schedules** tab, configure all the necessary schedules. For example, Full and incremental schedules.
- 5 Depending on the policy type that you selected, add the clients, database instances, or virtual machines that you want to protect. Perform this configuration on the **Clients** or the **Instances and databases** tab.
 - For most policy types you configure a list of clients on the **Clients** tab.
 - For **Oracle** and **MS-SQL-Server** policy types, you select instances or databases on the **Instances and databases** tab. Or if you use scripts or batch files, you select clients on the **Clients** tab.
- 6 Depending on the policy type that you selected, add the files, database instances, or other objects that you want to protect. This configuration is performed on the **Backup selections** tab.
- 7 For the policy types that have additional tabs, review and select the other policy options that are needed to complete the setup.
- 8 Click **Create**.

Example policy - Exchange Server DAG backup

This example describes how to create a policy to back up all databases in an Exchange Server DAG.

To add a policy for an Exchange Server DAG backup

- 1 On the left, select **Protection > Policies**.
- 2 Click **Add**.

- 3 On the **Attributes** tab, select the following:
 - **Policy type:** MS-Exchange-Server
 - **Perform snapshot backups:** Must be enabled.
 - **Enable granular recovery:** Optional. Enable this option if you want to restore individual mailbox and public folder objects from the database backup.
 - **Database backup source:** Choose whether to back up the active or the passive copy of the database. Also configure the preferred list, depending on the backup source that you selected.
- 4 On the **Schedules** tab, configure all the necessary schedules. For example, Full and incremental schedules.

Name	Type	Frequency	Retention
Full-backup	Full backup	1 week	2 weeks
Incremental-backup	Differential incremental	1 day	2 weeks

- 5 On the **Clients** tab, add one or more DAG names.

Client name	Hardware	Operating system
dag1234.domain.com	Windows-x64	Windows2016
dag5678.domain.com	Windows-x64	Windows2016

- 6 On the **Backup selections** tab, add the following directive.

```
Microsoft Exchange Database Availability Groups:\
```

Backup selection list

```
Microsoft Exchange Database Availability Groups:\
```

- 7 Click **Create**.

Example policy - Sharded MongoDB cluster

This example describes how to create a policy to back up the primary configuration server in a Sharded MongoDB cluster.

To add a policy for a MongoDB cluster backup

- 1
- On the left, select **Protection > Policies**.
- 2
- Click **Add**.
- 3
- On the **Attributes** tab, select the following:

■ **Policy type:** BigData
- 4
- On the **Schedules** tab, configure all the necessary schedules. For example, Full and incremental schedules.

Name	Type	Frequency	Retention
Full-backup	Full backup	1 week	2 weeks
Incremental-backup	Differential incremental backup	1 day	2 weeks

- 5
- On the **Clients** tab, add the client name. Use the format `MongoDBNode-portnumber`.

The following list backs up the primary configuration server on port 1.

Client name	Hardware	Operating system
primaryconfigserver-01	Linux	Red Hat 2.6.32

- 6
- On the **Backup selections** tab, add the application type, the backup hosts, and manually add the ALL_DATABASES directive.

Backup selection list	Notes
Application_Type=mongodb	The parameter values are case-sensitive.
mongodbhost=mongodbhost.domain.com	Use the format <code>Backup_Host=<FQDN_or_hostname></code> . The backup host can be a NetBackup client or a media server.
ALL_DATABASES	

- 7
- Click **Create**.

Example policy - Epic-Large-File

This example describes how to create a policy to back up very large database files such as the EPiC Database.

To add a policy for a large file

- 1 On the left, select **Protection > Policies**.
- 2 Select **Add**.
- 3 Go to the **Attributes** tab.
- 4 For the **Policy type** select **Epic-Large-File**.
- 5 On the **Schedules** tab, configure the full backup schedule.
- 6 On the **Clients** tab, add the client name.

Client name	Hardware	Operating system
primary1234.domain.com	Linux	Linux Red Hat 7.9, 8.x, or 9.x
primary5678.domain.com	Linux	SUSE 12 SP5+
primary1212.domain.com	AIX	AIX 7.2

- 7 On the **Backup selections** tab, add the path names.
- 8 On the **Epic-Large-File** tab, configure the following:

Number of parallel streams The number of parallel backup streams that are used for a backup selection.

If you have configured multiple backup selections in a policy, each backup selection has this number of streams. For example, assume that the number of streams per backup selection is 4 and there are two entries in the backup selections. In this case there are 4 concurrent streams for each backup selection with a total of 8 streams.

Use multiple MSDP nodes The option **Use multiple MSDP nodes** allows backups of data files in parallel, which accelerates the backup and restore jobs. You can use multiple MSDP storage units, a storage unit of an MSDP cluster, or a storage unit of an MVG (MSDP volume group) volume.

The dropdown list contains the available storage items, including MSDP storage units and SLPs that are configured in the system. Select the storage to which you want to distribute the data:

- MSDP storage units or SLPs. Select each storage unit that you want to use. If replication is needed, all SLP targets must be in the same target domain. (These targets include the policy storage SLP and the selected storage SLPs.) This configuration ensures that the data files are in the same domain after replication.
- A storage unit of an MSDP cluster.

Note: Select only this one storage unit. Do not select any other storage units with this selection.

This option distributes data files to multiple nodes in the cluster. The recommended value for **Number of parallel streams** is a multiple of the number of cluster nodes. If there are 4 nodes in the cluster, the value is **4n**. For example: 4, 8, or 12.

- A storage unit of an MVG (MSDP volume group) volume.

Note: Select only this one storage unit. Do not select any other storage units with this selection.

This option distributes data files to the MSDP disk volumes in the MVG volume. The recommended value for **Number of parallel streams** is a multiple of the number of MSDP disk volumes. If there are 4 MSDP disk volumes in the MVG volume, the value is **4n**. For example: 4, 8, or 12.

9 Select **Create**.

Edit, copy, or delete a policy

You can make changes to a policy, copy a policy, or delete a policy that you no longer need.

For details on policy options, refer to the *NetBackup Administrator's Guide, Volume I* and to the appropriate workload or database guides.

Note: You must have the RBAC Administrator role or similar permissions to manage policies.

Edit a policy

You can make changes to policy attributes, schedules, clients, or backup selections.

To edit a policy

- 1 On the left, select **Protection > Policies**.
- 2 Select the policy that you want to change and click **Edit**.
- 3 Make the changes that you want, then click **Save**.

Copy a policy

You can copy a policy to save time creating new policies. This option is especially useful for the policies that contain many of the same policy attributes, schedules, clients, or backup selections.

To copy a policy

- 1 On the left, select **Protection > Policies**.
- 2 Select the policy that you want to copy and click **Copy policy**.
- 3 Provide a name for the policy and click **Copy**.

Delete a policy

You can delete a policy if you no longer need it. To maintain protection of the clients or hosts, add them to another policy before you delete the current policy.

To delete a policy

- 1 On the left, select **Protection > Policies**.
- 2 Select one or more policies and click **Delete > Delete**.

Deactivate or activate a policy

Active policies are available for NetBackup to use to schedule backups or allow user-backups.

You can also use the **Go into effect at** policy attribute to activate a policy or to deactivate a policy. Or, to select a time for a policy to become active.

Deactivate a policy

You can deactivate a policy to temporarily pause any backup requests for that policy. For example, if you want to perform maintenance on the clients in the policy. Note that manual backups or user-requested backups cannot run if a policy is deactivated.

To deactivate a policy

- 1 On the left, click **Protection > Policies**.
- 2 Select the policy, then click **Deactivate**.

Activate a policy

Activate a policy when you are ready for backup schedules in the policy to run.

To activate a policy

- 1 On the left, click **Protection > Policies**.
- 2 Select the policy, then click **Activate**.

View automanaged policies and SLPs

You can filter the list of policies to also display automanaged policies and SLPs. NetBackup creates this type of policy when a workload administrator subscribes an asset to a protection plan. You cannot modify or delete these types of policies. You can only view the details.

To view automanaged policies and SLPs

- 1 On the left, select **Protection > Policies**.
- 2 In the toolbar, select the **Filter** icon.
- 3 Go to **Protection type** and select **Protection plan policies**. You can also sort by **Protection type** by clicking on the **Protection type** column.

The policy name for this type is made up of the protection plan name and a suffix that contains the GUID. If you used the custom prefix option when you created the protection plan, then the name of the automanaged policy reflects that prefix instead. For example:

```
vmware_backups+00000aa1-1aab-1a00-000f-0a0a00000a0b
```

- 4 To view the details of a policy, select the link for that policy.
The policy settings are displayed but grayed out.
- 5 Select **Close** to close the policy.

About automanaged policies or storage lifecycle policies

When a workload administrator subscribes an asset to a protection plan, NetBackup also creates an automanaged policy and automanaged storage lifecycle policy (SLP). In the NetBackup web UI, the name of an automanaged policy or SLP is made up of the protection plan name, then the suffix that contains the GUID. (If you used the custom prefix option when you created the protection plan, then the name of the automanaged policy reflects that prefix instead.) In the NetBackup

Administration Console, an automanaged policy or SLP name uses the prefix SLO_ENGINE_MANAGED+.

The NetBackup web UI does not allow users to modify or delete automanaged policies or SLPs. You can only view the details of an automanaged policy or SLP.

Warning: It is not recommended that users modify or delete automanaged policies or storage lifecycle policies using the **NetBackup Administration Console** or the command line. If a user begins to modify or delete an automanaged policy or SLP using the **NetBackup Administration Console**, a dialog appears that warns users about the possible consequences.

- If the user continues to make modifications, they must make sure that the policy or SLP continues to meet the service level objective that is the protection plan defines.
- If the user continues to delete the policy or SLP, they must make sure that the asset is added to another protection plan that meets the service level objective.

Perform manual backups

A manual backup is user-initiated and is based on a policy. For example, you can use a manual backup to prepare for the upcoming events that are outside scheduled backups, such as system maintenance.

In some cases, it may be useful to create a policy and schedule that is used only for manual backups. Create a policy for manual backups by creating a policy with a single schedule that has no backup window. Without a backup window, the policy can never run automatically.

To perform a manual backup

- 1 On the left, select **Protection > Policies**.
- 2 Select one or more policy names and click **Manual backup**.

To perform a manual backup, you must enable the **Go into effect at** attribute for the policy. If this attribute is set for a future date and time, the backup does not run.
- 3 Choose from the following options:

For a single policy:

 - Select the schedule and then clients that you want to back up.
 - Click **Backup**.

For multiple policies:

- To back up all the clients and the default schedules for the selected policies, click **Backup all**.
- To select specific clients and the schedule for each policy, click **Specify**.
- Follow the prompts to continue.

About the Epic-Large-File policy type

The Epic-Large-File policy uses multistreaming to speed up backup and restore performance for large files. The policy is targeted for the applications such as the medical record application EPiC Database for a single large file or a few large files.

Things to consider:

- An Epic-Large-File policy supports the MSDP storage units and the AdvancedDisk storage units. Some options in policy settings may apply only to MSDP storage units. For example, client-side deduplication applies only to MSDP storage units.

We recommend that you use the same storage type with the similar capabilities such as performance, capacity, and network bandwidth. If the storage is WORM, all of them must have the same retention settings.

- Settings to allow the parallel streams:
On the primary server, set **Global settings > Maximum jobs per client** to the appropriate value.
On the storage unit, set **Maximum concurrent jobs** to the appropriate value.
- For the Epic-Large-File policy, one backup selection may be split into multiple jobs. Each child job displays one file path under the **File List** in the Activity monitor.
- Create `bpstart_notify` and `bpend_notify` scripts for an Epic-Large-File policy. An Epic-Large-File policy ignores the generic `bpstart_notify` and `bpend_notify` scripts. You must include the `.<policyname>` or `.<policyname.schedule>` suffix to the script name or it does not run at the start or end of the policy.

Examples:

- UNIX

```
/usr/opensv/netbackup/bin/bpstart_notify.epic_file
/usr/opensv/netbackup/bin/bpend_notify.epic_file.full
```

- Windows

```
<installation_directory>\NetBackup\bin\bpstart_notify.epic_file.bat
<installation_directory>\bin\bpend_notify.epic_file.full.bat
```

For more information about the scripts, see the *NetBackup Administrator's Guide, Volume II*.

To restore the Epic-Large-File policy backups, use `nbepicfile` command. For more information, see the `nbepicfile` command in the *NetBackup Commands Reference Guide*.

See [“Example policy - Epic-Large-File”](#) on page 361.

Managing protection plans

This chapter includes the following topics:

- [Create a protection plan](#)
- [Customizing protection plans](#)
- [Edit or delete a protection plan](#)
- [Subscribe an asset or an asset group to a protection plan](#)
- [Unsubscribe an asset from a protection plan](#)
- [View protection plan overrides](#)
- [Copy a protection plan policy \(automanaged policy\) to a classic policy](#)
- [About Backup now](#)

Create a protection plan

Note: After upgrade, the protection plans may not appear in the web UI. The conversion process may not have run but should run within 5 minutes of performing the upgrade.

Before you create a protection plan, you must configure all storage options.

See [“About storage configuration”](#) on page 241.

To create a protection plan

- 1 On the left, click **Protection** > **Protection plans** and then click **Add**.
- 2 In **Basic properties**, enter a **Name**, **Description**, and select a **Workload** from the drop-down list.

Optional selection:

- **Policy name prefix:**
Use this option for policy names. A prefix is added to the policy name when NetBackup automatically creates a policy when users subscribe assets to this protection plan.
- **Enable Continuous Data Protection**
For VMware workloads, select this option to use Continuous data protection for the workload. Select the **Use universal share** option to use universal share for data storage. Using universal share substantially reduces the staging data storage requirements, thereby greatly reducing the data storage costs. CDP with universal share is supported NetBackup version 10.2 onwards. See the *Continuous data protection* chapter of the *NetBackup for VMware Administrator's Guide* for details.
- **Protect PaaS assets only**
For Cloud workloads, you must select this option to protect non-RDS PaaS assets with non-snapshot based protection, using the protection plan. Do not select this option for RDS assets with snapshot-based protection. See the *Managing PaaS assets* chapter of the *NetBackup Web UI Cloud Administrator's Guide* for details.

3 In **Schedules**, click **Add**.

If you have selected cloud as workload of Azure or Azure Stack, see the *Configuring backup schedules for cloud workloads* section of the [NetBackup Web UI Cloud Administrator's Guide](#).

You can set up a daily, weekly, or monthly backup and then set retention and replication of that backup. Also depending on workload, you can set up the following backup schedules: a **Automatic**, **Full**, **Differential incremental**, **Cumulative Incremental**, or **Snapshot only**.

For more information about AWS snapshot replication, review the *Configure AWS snapshot replication* in the [NetBackup Web UI Cloud Administrator's Guide](#)

If you select **Monthly** as a frequency, you can select between **Days of the week** (grid view) or **Days of the month** (calendar view).

Note: If you select **Automatic** for the schedule type, then all schedules for this protection plan are **Automatic**. If you select a **Full**, **Differential incremental**, or **Cumulative Incremental** for the schedule type, then all schedules for this protection plan must be one of these options.

If you select **Automatic** for the schedule type, NetBackup automatically sets the schedule type for you. NetBackup calculates when to do a **Full** or **Differential incremental** based on frequency you specify.

Note: The protection plan creation does not work for the VMware workload when certain schedule frequencies are set with WORM storage lock duration. The protection plan creation does not work when: schedule frequencies are set to less than one week and WORM storage **Lock Maximum Duration** less than one week greater than the requested retention period.

If you use a protection plan to protect VMware with WORM capable storage, set the WORM storage **Lock Maximum Duration** to greater than one week. Or, explicitly select the schedule type in the protection plan.

In the **Attributes** tab:

- Select the **Backup type**, how often it runs, and how long to keep the backup for this schedule.
 - The selections in the **Backup type** are dependent on workload that is selected and any other backup schedules that are currently active in this protection plan.
- (Optional) To replicate the backup, select **Replicate this backup**.
 - To use the **Replicate this backup** option, the backup storage must be a source in a targeted A.I.R. environment. The **Replication target** is configured in step 4.
 - For more information about replication, review *About NetBackup Auto Image Replication* in the [NetBackup Administrator's Guide, Volume I](#).
- (Optional) To keep a copy in long-term storage, turn on **Duplicate a copy immediately to long-term retention**. This option is not available for all workloads.
 - NetBackup immediately duplicates a copy to long-term storage after the backup completes.
 - The schedule options that are available for long-term storage are based on the frequency and the retention levels for the regular backup schedules that you created.

In the **Start window** tab:

- Define a **Start day**, **Start time**, **End day**, and **End time** for this schedule using the options available on the screen. Or you can drag your cursor over the time boxes to create the schedule.
- Use the options on the right to duplicate, remove, or undo changes to a schedule.

Click **Save** after all options are selected in the **Attributes** and the **Start window** tabs.

Review the **Backup schedule preview** window and verify that all schedules are set correctly.

- 4 In **Storage options**, configure the storage type per schedule you configured in step 3.

The options vary depending on storage options currently setup to work with NetBackup.

A protection plan can only use the storage that a NetBackup 8.1.2 or newer media server can access.

Storage option	Requirements	Description
Snapshot storage only	Snapshot Manager is required for this option.	
Perform snapshot backups	Microsoft SQL Server is required for this option.	For instructions on configuring protection plans for Microsoft SQL Server, see the <i>NetBackup Microsoft SQL Server Administrator's Guide</i> .
Backup storage	OpenStorage is required for this option. Tape, storage unit groups, and Replication Director not supported.	Click Edit to select the storage target. Click Use selected storage after selecting the storage target. The NetBackup Accelerator feature allows protection plans to run faster than traditional backups, by creating a compact data stream that uses less network bandwidth. If the storage server on the NetBackup primary server supports NetBackup Accelerator, this feature is included in the protection plan. For more details on NetBackup Accelerator, contact the NetBackup administrator or see the NetBackup Administrator's Guide, Volume I or the NetBackup for VMware Administrator's Guide. The Instant access feature allows the plan's recovery points to support the creation of instant access VMs or databases.
Replication target	The backup storage must be a source in a targeted A.I.R. environment.	Click Edit to select the replication target primary server. Select a primary server and then select a storage lifecycle policy. Click Use selected replication target to return to the storage options screen. Cloud workloads support the MSDP and MSDP-C storage units for replication (A.I.R.). If the replication target primary server is not in the list, you must add one in NetBackup. For more information on how to add a replication target primary server, review <i>Adding a trusted primary server</i> in the NetBackup Deduplication Guide.

Storage option	Requirements	Description
Long-term retention storage	OpenStorage is required for this option. Tape, storage unit groups, and Replication Director not supported.	Click Edit to select the cloud storage provider. Click Use selected storage after selecting the cloud provider target. Cloud workloads support the AdvancedDisk, Cloud storage, MSDP, and MSDP-C as storage units for duplication.
Transaction log options	Microsoft SQL Server is required for this option.	If you use the option Select custom storage options , click Edit to select the backup storage.

5 In **Backup options**, configure all options based on your workload type. The options in this area change depending on workload, schedule, or storage options selected.

For the **Cloud** workload:

- For any of the selected cloud provider options, if you select **Enable granular recovery for files or folders**, ensure that you have opted to retain a snapshot while adding a backup schedule, as granular recovery can be performed only from a snapshot image.
- For any of the selected cloud provider option, if you select **Exclude selected disks from backups**, then the selected disks would not be backed-up and hence the VM would not be recovered completely. Any application running on the excluded disks might not work.

Note: The boot disks cannot be excluded from the backups even if they have data or tags associated with them.

- If you have selected the cloud provider as Google Cloud Platform, select **Enable regional snapshot**, to enable regional snapshots.
If the regional snapshot option is enabled, the snapshot is created in the same region in which the asset exists. Otherwise, the snapshot is created in a multi-regional location.
- (Microsoft Azure or Azure Stack Hub cloud provider) Select **Specify snapshot destination resource group** to associate snapshots to a particular peer resource group. This resource group is within the same

region in which the asset exists. Select a configuration, subscription, and a resource group for a snapshot destination.

- If you have selected **Enable Continuous data protection** for a VMware workload, select a Continuous data protection gateway from the list. Click **Next**. If you are using the universal share option, the gateway version must be NetBackup 10.2 or higher.

6 In **Permissions**, review the roles that have access to protection plans.

To give another role access to this protection plan, click **Add**. Select the **Role** in the table and customize the role by adding or removing permissions in the **Select permissions** section.

7 In **Review**, verify that the protection plan details are correct and click **Finish**.

Customizing protection plans

After you create a protection plan, only certain settings are available to change or configure. See [Table 21-1](#).

Table 21-1 Protection plan settings that can be configured and edited

Protection plan setting	Setting is available when you...		Notes
	Edit a plan	Subscribe an asset	
Storage options	X		
Backup options		X	
Advanced options		X	
Schedules	X	X	Backup window only. For SQL Server, transaction log frequency, and retention.
Protected assets		N/A	
Permissions	X	N/A	Can add a role.

Edit or delete a protection plan

Edit a protection plan

You can make changes to the **Description**, **Storage options**, and **Schedules** (limited) for a protection plan.

Note: You cannot edit these settings in a protection plan: **Backup options** and **Advanced options**. If you want to adjust these settings and additional schedule settings, you must create a new protection plan and subscribe assets to the new plan. Or, you can customize the plan for the asset.

See [“Customizing protection plans”](#) on page 374.

To edit a protection plan

- 1 On the left, click **Protection > Protection plans**.
- 2 Click on the protection plan name that you want to edit.
- 3 Click **Edit description** to edit the description.
- 4 (Optional) In the **Storage options** section, click **Edit** to change the storage options.

Delete a protection plan

You cannot delete a protection plan unless all assets have been removed from the protection plan. If you want to maintain protection on the assets, add another protection plan to those assets before you delete the current protection plan.

See [“Unsubscribe an asset from a protection plan”](#) on page 377.

See [“Subscribe an asset or an asset group to a protection plan”](#) on page 376.

See [“Create a protection plan”](#) on page 368.

To delete a protection plan

- 1 On the left, click **Protection > Protection plans**.
- 2 Select the check box for the protection plan that you want to delete.
- 3 Click **Delete > Yes**.

Subscribe an asset or an asset group to a protection plan

You can subscribe a single asset or a group of assets to a protection plan. An asset or a group of assets can be subscribed to multiple protection plans. Before you can subscribe assets to a protection plan, you must create a protection plan.

NetBackup supports homogenous cloud asset subscriptions. When you subscribe an asset to a protection plan, the cloud provider of the asset must be the same as the cloud provider that is defined in the protection plan.

Note: You cannot edit these settings when you subscribe an asset: **Storage options** or **Permissions**. Changes to **Schedules** are limited. If you want to adjust these settings, you must create a new protection plan and subscribe assets to the new plan. Or, you can customize the plan for the asset.

See [“Customizing protection plans”](#) on page 374.

To subscribe an asset or an asset group to a protection plan

- 1 On the left, click **Workloads** then the workload type (for example: **VMware**).
- 2 Select an asset type (for example: **Virtual machines**, **Intelligent VM groups**).
- 3 Select one or more assets.
- 4 Click **Add protection**.

If you selected a Cloud workload asset or asset group, proceed to step 7.

- 5 In **Choose a protection plan**, select the name of the protection plan and click **Next**.
- 6 (Optional) Adjust any options in the **Backup options** or **Advanced options**.

- **Schedules**

Change the backup start window for full or incremental schedules.

For SQL Server transaction log schedules you can change the start window, the recurrence, and the retention period.

- **Backup options**

Adjust the backup options that were set up in the original protection plan. The options in this area change depending on workload.

- **Advanced**

Change or add any options that were set up in the original protection plan.

You need the following permissions to make these changes:

- **Edit attributes**, to edit **Backup options** and **Advanced options**.

- **Edit full and incremental schedules**, to edit the start window for these schedule types.
- **Edit transaction log schedules**, to edit the settings for SQL Server transaction log schedules.

7 Click **Protect**.

Unsubscribe an asset from a protection plan

You can unsubscribe individual assets or groups of assets from a protection plan.

Note: When you unsubscribe an asset from a protection plan, there is a possibility that the asset displays **Classic policy** in the web UI. This situation can happen when an asset is subscribed to a protection plan and a backup runs for that asset. Then the asset is unsubscribed from the protection plan while it has a valid backup image. The web UI displays **Classic policy**, but there may or may not be an active policy protecting the asset.

To unsubscribe a single asset from a protection plan

- 1 On the left, click **Workloads** then the workload type (for example: **VMware**).
- 2 Select a single asset type (for example: **Virtual machines**).
- 3 Click on the specific asset name.
- 4 Click **Remove protection** and click **Yes**.

To unsubscribe a group of assets from the protection plan

- 1 On the left, click **Workloads** then the workload type (for example: **VMware**).
- 2 Select a group asset type (for example: **Intelligent VM groups**).
- 3 Click on the specific group asset name.
- 4 Click **Remove protection** and click **Yes**.

View protection plan overrides

When you set permissions for protection plans, you can set the permissions to allow your workload administrator to customize assets a protection plan covers. The workload administrator can apply overrides to certain areas of schedules and backup options for an asset.

To view protection plan overrides

- 1 On the left, click **Protection > Protection plans** and then click the name of the protection plan.
- 2 In the **Protected assets** tab, click on **Applied** in the **Custom settings** column.
- 3 Review the original and the new settings in the **Schedules** and **Backup options** tabs.
 - **Original:** The setting when the protection plan was first created.
 - **New:** The last change that was made to the protection plan for that setting.

Copy a protection plan policy (automanaged policy) to a classic policy

When a workload administrator subscribes an asset to a protection plan, NetBackup also creates an automanaged policy and an automanaged storage lifecycle policy (SLP). You cannot modify or delete automanaged policies or SLPs. You can only view the details of this type of policy.

If you prefer to use a classic policy for backups, you can copy the settings in an automanaged policy to a classic policy. This action is only available when a classic policy type exists that corresponds to the workload. For other protection plan types, you cannot copy the automanaged policy to a classic policy.

To copy an automanaged policy to a classic policy

- 1 On the left, select **Protection > Policies**.
- 2 To locate the automanaged policy, in the toolbar select the **Filter** icon.
- 3 Go to **Protection type** and select **Protection plan policies**. You can also sort by **Protection type** by clicking on the **Protection type** column.

The policy name for this type is made up of the protection plan name and a suffix that contains the GUID. If you used the custom prefix option when you created the protection plan, then the name of the automanaged policy reflects that prefix instead. For example:

```
vmware_backups+00000aa1-1aab-1a00-000f-0a0a00000a0b
```

- 4 Select **Actions > Copy policy**.
- 5 Enter a new policy name. Then select **Copy**.

NetBackup validates the settings in the policy and opens the policy.

- 6 After you copy an automanaged policy, the new classic policy is configured to use an automanaged SLP. Change the storage for the policy to a storage that you manage.
 - On the **Attributes** tab, note that a value in the **Policy storage** is set to **Any available**.
 - Go to the **Schedules** tab.
 - Locate and select the schedule. Then select **Edit**.
 - Go to the **Attributes** tab.
- 7 You have the following choices for the storage selection.

Choose not to override the policy storage with the storage that is configured for the schedule. Instead, select the storage on the policy **Attributes** tab.

To configure storage for the policy:

- Clear the check box **Override policy storage selection**. Then select **Save**.
- Go to the policy **Attributes** tab.
- Locate the **Policy storage** list and then select the storage that you want to use.

In the schedule attributes, select a different storage that is not auto-managed.

To configure storage for a schedule:

- Locate **Override policy storage selection**. Note that the storage that is selected in the list is an automanaged storage.
- Select a storage that you manage.
- Select **Save**.

- 8 Review the other policy settings and make any wanted changes.
- 9 Select **Save**.

About Backup now

With Backup now, workload administrators can back up an asset immediately. For example, you can use Backup now to prepare for the upcoming events that are outside scheduled backups, such as system maintenance. This type of backup is independent of scheduled backups and does not affect future backups. You can manage and monitor a Backup now job in the same way you manage and monitor other NetBackup jobs. Note that Backup now jobs cannot be restarted.

Backup now is supported for the following workloads:

- Cassandra
- Cloud and PaaS

NetBackup supports homogenous cloud asset subscriptions. While you subscribe an asset to a protection plan, the cloud provider of the asset must be the same as the cloud provider defined in the protection plan.

- Kubernetes
- Microsoft SQL Server
- MySQL
- Nutanix AHV
- PostgreSQL
- RHV
- VMware

Note: To use Backup now you must have subscribe permissions for at least one protection plan. You can select only one asset at a time for each Backup now operation.

Immediately back up an asset using Backup now

You can start Backup now for an asset from the list of assets. For example, from the list of virtual machines, intelligent groups, or databases. Or, you can start Backup now from the asset's details. These details display all of the protection plans to which the asset is subscribed. You can choose **Backup now** from any one of these protection plans.

To immediately back up an asset using Backup now

- 1 On the left, select the workload and locate the asset that you want to back up.
- 2 Select **Actions > Backup now**.

3 Choose a protection plan for the backup.

All protection plans to which the asset is subscribed are listed.

To back up an asset that is not subscribed to any protection plan, select **Backup now** and choose from the existing protection plans. You can also create a new protection plan and then use it with a **Backup now** operation.

Note: The option of **Backup type** is only available for Microsoft SQL Server assets. You can select the type of backup you want to perform using the drop-down. The drop-down only contains the backup types that are available in the protection plan.

4 Click **Start backup**.

Protecting the NetBackup catalog

This chapter includes the following topics:

- [About the NetBackup catalog](#)
- [Catalog backups](#)
- [Disaster recovery emails and the disaster recovery files](#)
- [Disaster recovery packages](#)
- [Set the passphrase to encrypt disaster recovery packages](#)
- [Recovering the catalog](#)

About the NetBackup catalog

A NetBackup catalog is the internal database that contains information about NetBackup backups and configuration. Backup information includes records of the files that have been backed up and the media on which the files are stored. The catalogs also contain information about the media and the storage devices.

Configure a disaster recovery pass phrase and a catalog backup before you run any regular backups. NetBackup needs information from the catalog to determine where the backups of files are located. Without a catalog, NetBackup cannot restore data.

See [“Set the passphrase to encrypt disaster recovery packages”](#) on page 391.

See [“Configuring catalog backups”](#) on page 385.

As additional protection for the catalog, consider archiving the catalog.

See [“Archiving the catalog and restoring from the catalog archive”](#) on page 657.

Catalog backups

Because the catalog plays an integral part in a NetBackup environment, a special type of backup protects the catalog and is separate from regular client backups. A catalog backup policy backs up catalog-specific data as well as produces disaster recovery information. The catalog can be stored on a variety of media.

The catalog backup is designed for active environments in which continual backup activity occurs. It includes all the necessary catalog files, the databases (NBDB, NBAZDB, and BMRDB), and any catalog configuration files. The catalog backup can be performed while regular backup activity occurs. Incremental backups of a large catalog can significantly reduce backup times.

Configure a catalog backup before you run any regular backups. NetBackup needs information from the catalog to determine where the backups of files are located. Without a catalog, NetBackup cannot restore data.

See [“Configuring catalog backups”](#) on page 385.

As additional protection for the catalog, consider archiving the catalog.

See [“Archiving the catalog and restoring from the catalog archive”](#) on page 657.

From a catalog backup an administrator can recover either the entire catalog or pieces of the catalog. (For example, separately recover the databases from the configuration files.) Details about catalog recovery scenarios and procedures are available in the *NetBackup Troubleshooting Guide*.

The catalog backup process

The catalog backup performs the following tasks:

- Backs up the catalog while continual client backups are in progress.
- Performs a full or an incremental catalog backup.
- Runs the scheduled catalog backups.
- Copies the databases to the staging directory and then backs up that directory.
- Creates the disaster recovery package.
- Catalog backups that are made to tape also include the following items:
 - Spans multiple tapes for a catalog backup.
 - Allows for a flexible pool of catalog tapes.
Catalog backups to tape use media from the **CatalogBackup** volume pool only.
 - Appends to existing data on tape.

- When an online catalog backup is run, it generates three jobs: A parent job, a child job for NetBackup relational database tables, and a child job for catalog images and configuration data. The child jobs contain the actual backed up data. Consider both child jobs to duplicate, verify, or expire the backup.

Refer to the following topics for information on how to configure a catalog backup:

See [“Prerequisites for backing up the NetBackup catalog”](#) on page 384.

See [“Configuring catalog backups”](#) on page 385.

Prerequisites for backing up the NetBackup catalog

The following prerequisites exist for a catalog backup:

- Set a passphrase for the disaster recovery package.
See [“Disaster recovery packages”](#) on page 390.
See [“Set the passphrase to encrypt disaster recovery packages”](#) on page 391.
If the passphrase is not set, catalog backups fail.
- The primary server and the media server must both be at the same NetBackup version.
See the [NetBackup Installation Guide](#) for information about mixed version support.
- Catalog backups write only to media in the **CatalogBackup** volume pool. A storage device must be configured and media must be available in the **CatalogBackup** volume pool.
- The following requirements exist if the primary server is configured to use a non-privileged user (or service user) account. For more information on this type of account, refer to the [NetBackup Security and Encryption Guide](#).
 - The service user account must have the write access permissions on the disaster recovery (DR) path.
 - Configure the catalog policy with the credentials for the service account. (This setting is located on the **Disaster recovery** tab.)
 - You cannot use another user account, even if that account has the access to the DR path. The NetBackup Administrator must ensure that the service user can write to any network share without switching the context to another user.
On Windows, this requirement is not applicable if the DR path is a network share.

Configuring catalog backups

To protect the NetBackup catalog, you create a backup policy that is specific for catalog backups.

For information on how to configure catalog backups in Windows clustered environments, see the [NetBackup Clustered Primary Server Administrator's Guide](#).

To configure a catalog backup

- 1 Review the prerequisites for performing catalog backups.
See [“Prerequisites for backing up the NetBackup catalog”](#) on page 384.
- 2 Sign in to the NetBackup web UI.
- 3 Click **Protection > Policies**. Then click **Add**.
- 4 On the **Attributes** tab, complete the following entries:
 - Enter a unique **Policy name**.
 - For the **Policy type**, select **NBU-Catalog**.
 - **Policy storage**
For disk storage units, increase the **Maximum Concurrent Jobs** storage unit setting to ensure that the catalog backup can proceed during regular backup activity.

Note: If your installation contains media servers at various versions, you can select a specific media server for the destination **Policy storage**. Do not select **Any Available**.

- **Policy volume pool**
NetBackup automatically creates a **CatalogBackup** volume pool that is selected by default only for **NBU-Catalog** policy types.
- For other policy attribute descriptions, see the following topic:
- 5 On the **Schedules** tab, configure the schedules you want for the catalog backup.
See [“Concurrently running catalog backups with other backups”](#) on page 387.
See [“Catalog policy schedule considerations”](#) on page 387.
- 6 Click the **Disaster recovery** tab.
The tab contains information regarding the location of data crucial to disaster recovery.

- Provide the path where each disaster recovery image file can be saved on disk. Enter the **Network share username** and **Network share password**, if necessary.
It is recommended that you use a network share or a removable device.
Do not save the disaster recovery information to the local computer.
- 7 Select **Send disaster recovery email** and enter one or more email addresses for NetBackup administrators (separated by commas).

After every catalog backup, NetBackup sends disaster recovery information to the administrators that are indicated here.

Make sure that email notification is enabled in your environment.

See [“Disaster recovery emails and the disaster recovery files”](#) on page 389.
- 8 Add the policies that back up any critical data to the **Critical policies** list.

These policies are any that you consider crucial to the recovery of a site in the event of a disaster. The disaster recovery report includes a list of the media that is used for backups of critical policies. The report includes media only for incremental and full backup schedules, so any critical policies should use only incremental or full backup schedules.
- 9 Click **Create**.

Backing up NetBackup catalogs manually

Catalog backups typically run automatically per the **NBU-Catalog** policy. You can also manually start a catalog backup.

A manual catalog backup is useful in the following situations:

- To perform an emergency backup. For example, if the system is scheduled to be moved and you cannot wait for the next scheduled catalog backup.
- If there is only one standalone drive and the standalone drive is used for catalog backups. In this situation, automatic backups are not convenient. The catalog backup tape must be inserted before each catalog backup and removed when the backup is done. (The tape swap is necessary because NetBackup does not mix catalog and regular backups on the same tape.)

To perform a manual catalog backup

- 1 Sign in to the NetBackup web UI.
- 2 Click **Protection > Policies**.
- 3 Select the catalog backup policy that you want to run.
- 4 Click **Manual backup**.

- 5 (Optional) Select the schedule that you want to use.
- 6 Click **Backup**.

Concurrently running catalog backups with other backups

You can schedule catalog backups to run concurrently with other backup types for the primary server.

Make the following adjustments to ensure that the catalog backup can proceed while regular backup activity occurs:

- Set the **Maximum jobs per client** value to greater than one. The property is found in the Global attributes host properties for the primary server.
- Increase the **Maximum concurrent jobs** setting on the storage unit where the backups are sent.

See [“Determining whether or not a catalog backup succeeded”](#) on page 388.

See [“Strategies that ensure successful NetBackup catalog backups”](#) on page 389.

Catalog policy schedule considerations

When you work with catalog policy schedules, consider the following:

- Schedule the catalog backups to occur on a regular basis. Without regular catalog backups, you risk losing regular backups if there is a problem with the disk that contains the catalogs.
- The following backup types are supported:
 - Full
 - Differential incremental
This incremental schedule depends on a full schedule.
 - Cumulative incremental
- The least frequent schedule runs if many schedules are due at the same time.
- One catalog backup policy can contain multiple incremental schedules that are session-based:
 - If one is cumulative and the others are differential, the cumulative runs when the backup session ends.
 - If all are cumulative or all are differential, the first schedule that is found runs when the backup session ends.
- The queued scheduled catalog backup is skipped if a catalog backup job from the same policy is running.

- Session end means that no jobs are running. (This calculation does not include catalog backup jobs.)
- The Vault catalog backup is run whenever triggered from Vault, regardless of whether a catalog backup job is running from the same policy.

How catalog incrementals and standard backups interact on UNIX

A catalog backup policy can include both full catalog backups and incremental catalog backups. However, incremental catalog backups differ from incremental standard backups. Catalog backups use both `mtime` and `ctime` to identify changed data. Standard incremental backups use only `mtime` to identify changed data.

Because of this difference, running a standard policy type backup that includes the `/usr/openv/netbackup/db/images/` directory can adversely affect incremental catalog backups. When standard backups run, they reset the file access time (`atime`). In turn, the reset changes the `ctime` for files and directories. If an incremental catalog backup runs, it sees that the `ctime` has changed and backs up the files. The backup may be unnecessary since the files may not have changed since the last catalog backup.

To avoid additional processing during catalog backups, the following is recommended:

If incremental catalog backups are configured, exclude the NetBackup `/usr/openv/netbackup/db/images/` directory from standard backups.

To exclude that directory, create a `/usr/openv/netbackup/exclude_list` file on the primary server.

See [“About NetBackup primary server installed directories and files”](#) on page 682.

Determining whether or not a catalog backup succeeded

The All Log Entries, Problems, and Media Log reports, available from the Reports utility, provide information on NetBackup catalog backups. The All logs entries report is also available in the NetBackup web UI.

An email message is sent to the address that is indicated in the **Disaster recovery** settings for a catalog backup.

Configure this email with the `mail_dr_info.cmd` (on Windows) or the `mail_dr_info` script (on UNIX).

See the [Administrator's Guide, Volume II](#) for more information on setting up this script.

Strategies that ensure successful NetBackup catalog backups

Use the following strategies to ensure successful catalog backups:

- Use only the methods that are described in this chapter to back up the catalogs. These are the only methods that can track all relevant NetBackup activities and ensure consistency between the catalog files.
- Back up the catalogs often. If catalog backup files are lost, the changes that were made between the last catalog backup and the time of the disk crash are lost.
- If you back up your catalogs to disk, always back up to a different disk than where the catalog files reside. If you back up the catalog to the disk where the actual catalog resides, both catalog backups are lost if the backup disk fails. Recovering the catalog is much more difficult. Also, ensure that the disk has enough space for the catalogs. Backups to a full disk fail.

Note: If a catalog backup is on tape, the tape must be removed when the backup is finished or regular backups cannot proceed. NetBackup does not mix catalog and regular backups on the same tape.

Disaster recovery emails and the disaster recovery files

In a catalog backup policy, you can configure the policy to send the disaster recovery information to an email address. This information appears on the **Disaster recovery** tab.

The disaster recovery email and the accompanying attachments that are sent contain the following important items for a successful catalog recovery:

- A list of the media that contains the catalog backup.
- A list of critical policies.
- Instructions for recovering the catalog.
- The image file as an attachment.

If a catalog backup policy included both full backups and incremental backups, the attached image file can be a full or an incremental catalog backup.

Recovering from an incremental catalog backup completely recovers the entire catalog if the **Automatically recover the entire NetBackup catalog** option is selected on the wizard panel. The entire catalog is recovered because the incremental catalog backup references information from the last full backup.

You do not need to recover the last full catalog backup before you recover the subsequent incremental backups.

- The disaster recovery package (`.drpkg` file) as an attachment.

Note: If you are not able to receive the disaster recovery packages over emails even after the disaster recovery email configuration, and then ensure the following:

Your email exchange server is configured to have the attachment size equal to or greater than the disaster recovery package size. You can check the size of the package (`.drpkg` file size) on the disaster recovery file location that you have specified in the catalog backup policy.

The firewall and the antivirus software in your environment allows the files with the `.drpkg` extension (which is the extension of a disaster recovery package file).

NetBackup emails the disaster recovery file when the following events occur:

- The catalog is backed up.
- A catalog backup is duplicated or replicated.
- The primary catalog backup or any copy expires automatically or is expired manually.

On Windows: You can tailor the disaster recovery email process by providing the `mail_dr_info.cmd` script in the `install_path\Veritas\NetBackup\bin` directory. This script is similar to the `nbmail.cmd` script. See the comments in the `nbmail.cmd` script for use instructions.

Disaster recovery packages

For increased security, a disaster recovery package is created during each catalog backup. The disaster recovery package file has `.drpkg` extension.

The disaster recovery (DR) package stores the identity of the primary server host. NetBackup requires this package to get the identity of the primary server back after a disaster. Once you have recovered the host identity, you can perform the catalog recovery.

The disaster recovery package contains the following information:

- NetBackup CA-signed certificates and private keys of the primary server certificate and the NetBackup certificate authority (CA) certificate

- Information about the hosts in the domain
- Security settings
- External CA-signed certificates
External CA-signed certificates from Windows certificate store, if applicable
- NetBackup configuration options that are specific to external CA-signed certificates
- Key management service (KMS) configuration

Note: By default, the KMS configuration is not backed up during catalog backup. Set the KMS_CONFIG_IN_CATALOG_BKUP configuration option to 1 to include the KMS configuration as part of the disaster recovery package during catalog backup.

Note: You must set a passphrase for the disaster recovery package for the catalog backups to be successful.

Set the passphrase to encrypt disaster recovery packages

During each catalog backup, a disaster recovery package is created and encrypted with the passphrase that you set. If you need to perform disaster recovery, you need to provide this encryption passphrase when you install NetBackup on the primary server in the disaster recovery mode.

If you do not set a passphrase before you run a catalog backup, the following points apply:

- NetBackup prevents you from configuring a new catalog backup policy.
- If the catalog backup policy is upgraded from a previous version, catalog backups continue to fail until the passphrase is set.

Note: Catalog backups may fail with status code 144 even though the passphrase is set. This situation occurs because the passphrase may be corrupted. To resolve this issue, you must reset the passphrase.

Caution: Ensure that the passphrase contains only the supported characters. If you enter a character that is not supported, you may face problems during disaster recovery package restore. The passphrase may not be validated and you may not be able to restore the disaster recovery package.

Set or modify the passphrase for disaster recovery packages (NetBackup web UI)

Before you modify the passphrase, review the following information:

See [the section called “Notes for modifying the passphrase for the disaster recovery packages”](#) on page 393.

To set or modify the passphrase (NetBackup web UI)

- 1 Open the NetBackup web UI.
- 2 At the top, select **Settings > Global security**.
- 3 Select **Disaster recovery**.
- 4 Enter and confirm the passphrase.

Review the following password rules:

- The existing passphrase and the new passphrase must be different.
- By default, the passphrase must contain a minimum of 8 and a maximum of 1024 characters.

You can set the passphrase constraints using the `nbseccmd -setpassphraseconstraints` command option.

- Only the following characters are supported for the passphrase: White spaces, uppercase characters (A to Z), lowercase characters (a to z), numbers (0 to 9), and special characters.
Special characters include: ~ ! @ # \$ % ^ & * () _ + - = ` { } [] ; ' , . / ? < > "

Caution: If you enter a character that is not supported, you may face problems during disaster recovery package restore. The passphrase may not be validated and you may not be able to restore the disaster recovery package.

- 5 Select **Save**. If the passphrase already exists, it is overwritten.

Set or modify the passphrase for disaster recovery packages (command-line interface)

Before you modify the passphrase, review the following information:

See [the section called “Notes for modifying the passphrase for the disaster recovery packages”](#) on page 393.

To set or modify the passphrase using the command-line interface

- 1 The NetBackup administrator must be logged on to the NetBackup Web Management Service to perform this task. Use the following command to log on:

```
bpnbat -login -loginType WEB
```

- 2 Run the following command to set a passphrase to encrypt disaster recovery packages:

```
nbseccmd -drpkgpassphrase
```

- 3 Enter the passphrase.

If a passphrase already exists, it is overwritten.

Notes for modifying the passphrase for the disaster recovery packages

Consider the following points before you modify the passphrase:

- Subsequent disaster recovery packages are encrypted with the new passphrase that you set.
- If you change the passphrase anytime, it is not changed for the previous disaster recovery packages. Only new disaster recovery packages are associated with the new passphrase.
- The passphrase that you provide when you install NetBackup on the primary server in a disaster recovery mode after a disaster must correspond to the disaster recovery package from which you want to recover the primary server host identity.

Recovering the catalog

Catalog recovery is discussed in the [NetBackup Troubleshooting Guide](#).

Managing backup images

This chapter includes the following topics:

- [About the Catalog utility](#)
- [Catalog utility search criteria and backup image details](#)
- [Verify backup images](#)
- [Promote a copy to a primary copy](#)
- [Duplicate backup images](#)
- [Expire backup images](#)
- [About importing backup images](#)

About the Catalog utility

Use the **Catalog** utility to search for a backup image to perform the following actions:

- Verify the backup contents with what is recorded in the NetBackup catalog.
See [“Verify backup images”](#) on page 397.
- Duplicate the backup image to create up to 10 copies.
- See [“Duplicate backup images”](#) on page 399.
- Promote a copy of a backup to be the primary backup copy.
- See [“Promote a copy to a primary copy”](#) on page 398.
- Expire backup images.
See [“Expire backup images”](#) on page 403.
- Import expired backup images or images from another NetBackup server.
See [“About importing expired images”](#) on page 404.

Catalog utility search criteria and backup image details

The catalog utility in the NetBackup web UI lets you perform various actions on a catalog image. For example, verify or duplicate an image. The catalog utility is organized as follows:

- **Search tab**
Provides the search criteria you can use to locate backup images. See [Table 23-1](#) for details.
For more details on these actions and on data-in-transit encryption (DTE) in your NetBackup environment, see the [NetBackup Administrator's Guide, Volume I](#) and [NetBackup Security and Encryption Guide](#).
After you search for backup images, the image list displays at the bottom of the page. Click **Show or hide columns** to display additional information about the images. See [Search results properties](#) for additional properties that are displayed in the search results.
- **Activity tab**
Displays the progress of the request to verify, duplicate, expire, or import an image.

Search criteria

The following actions and search criteria are available when you search for catalog images.

Table 23-1 Catalog search criteria

Property		Description
Action		Specifies the action that was used to create the image: Verify, Duplicate, Import . See "Verify backup images" on page 397. See "Duplicate backup images" on page 399. See "Expire backup images" on page 403.
Media		
	Media ID	The media ID for the volume. To search on all media, select <All> .
	Media host	The host name of the media server that produced the originals. To search all hosts, select All media hosts .
	Disk type	The disk type of the storage unit.
	Disk pool	The name of the disk pool. Not enabled if the disk type is BasicDisk.

Table 23-1 Catalog search criteria (*continued*)

Property		Description
	Media server	The name of the media server that produced the original images. To search all media servers, select All media hosts .
	Volume	The ID of the disk volume in the disk pool. Enabled if the disk type is not BasicDisk.
	Path	Searches for an image on a disk storage unit, if the path is entered. Or, searches all of the disk storage on the specified server, if All was selected. Enabled if the disk type is BasicDisk.
Date/time range		The range of dates and times that you want to search. The Global attributes property Policy update interval determines the default range.
Copies, policies, and clients		
	Copies	The copy that you want to search. Select either Primary or the copy number.
	Policy name	The policy under which the selected backups were performed. To search all policies, select All policies .
	Policy type	The purpose of the policy.
	Type of backup	The type of schedule that created the backup. To search all schedule types, select All backup types . Enabled if you select a specific Policy type .
	Client (host name)	The host name of the client that produced the backup. To search all hosts, select All clients .
Job priority		
	Override default job priority	The job priority for the catalog action (verify, duplicate, or import). To change the default, enable Override default priority. Then, select a value for the Job priority. If this option is not enabled, the job runs using the default priority as specified in the Default job priorities host property. Changes that you make affect the priority for the selected job only.
	Job priority	The priority of the catalog job. Enabled if you override the default priority.

Search results properties

In addition to properties that you can select for the search, other properties are displayed for the images.

Table 23-2 Catalog search results properties

Property	Description
Copy DTE mode	Specifies whether the data is transferred over a secure channel when the current image copy is created.
Copy hierarchy DTE mode	Specifies whether the data is transferred over a secure channel when the current image copy and all its parent copies in the hierarchy are created.
Expiration date	The date that the image expires.
Image DTE mode	Indicates the data-in-transit encryption (DTE) mode for the backup image.
Immutable	Indicates if the backup image is read-only and cannot be modified, corrupted, or encrypted.
Indelible	Indicates if the backup image is protected from being deleted before it expires.
Malware scan status	The scan status of the backup image.
Infection status	Displays the malware infected status of the backup images. The infection can be detected by malware scan or file hash search.
Mirror copy	Indicates if the image is a mirror replica or copy.
On hold	Indicates whether the image copy is on hold or not. Yes: The image has only one copy and a hold is set on the copy. No: No hold is set on the copy. A hold is set with the <code>nbholdutil</code> command.
Time	The time that the backup ran.
WORM unlock time	Indicates the time at which the image can be altered or deleted. Applies to the storage units that are WORM capable.

Verify backup images

NetBackup can verify the contents of a backup by reading the volume and comparing its contents to what is recorded in the NetBackup catalog.

This operation does not compare the data on the volume to the contents of the client disk. However, the operation does read each block in the image to verify that the volume is readable. (However, data corruption within a block is possible.) NetBackup verifies only one backup at a time and tries to minimize media mounts and positioning time.

To verify backup images

- 1 On the left, click **Catalog**.
- 2 From the **Action** list, select **Verify**.
- 3 Select the search criteria to find the image you want to verify. Click **Search**.
Backups that have fragments on another volume are included, as they exist in part on the specified volume.
See [“Catalog utility search criteria and backup image details”](#) on page 395.
- 4 Select the image that you want to verify. Then click **Verify**.
- 5 Click the **Activity** tab to view the job results.

Promote a copy to a primary copy

Each backup is assigned a primary copy. NetBackup uses the primary copy to satisfy restore requests. The first backup image that is created successfully by a NetBackup policy is the primary backup. If the primary copy is unavailable and a duplicate copy exists, select a copy of the backup and set it to be the primary copy.

NetBackup restores from the primary backup, and Vault duplicates from the primary backup. If your Vault profile performs duplication, you can designate one of the duplicates as the primary. In most circumstances, the copy remaining in the robot is the primary backup. When a primary backup expires, the next backup (if it exists) is promoted to primary automatically.

Use one of the following methods to promote a copy to a primary copy:

Promote a backup copy to a primary copy

See [the section called “Promote a backup copy to a primary copy”](#) on page 399.

Promote a copy to a primary copy for many backups using the `bpchangeprimary` command

See [the section called “Promote a copy to a primary copy for many backups”](#) on page 399.

Promote a backup copy to a primary copy

To promote a backup copy to a primary copy

- 1 On the left, click **Catalog**.
- 2 From the **Action** list, select **Duplicate**.
- 3 Select the search criteria to find the image you want to promote. Be sure that you indicate a copy in the **Copies** field and not **Primary copy**.
See “[Catalog utility search criteria and backup image details](#)” on page 395.
- 4 Click **Search**.
- 5 Select the image you want to promote. Then click **Set primary copy**.
After the image is promoted to the primary copy, the **Primary copy** column immediately reads **Yes**.
- 6 Click the **Activity** tab to view the job results.

Promote a copy to a primary copy for many backups

More information on the `bpchangeprimary` is available in the [NetBackup Commands Reference Guide](#).

To promote a copy to a primary copy for many backups

- ◆ You can also promote a copy to be a primary copy for many backups using the `bpchangeprimary` command. For example, the following command promotes all copies on the media that belongs to the `b_pool` volume pool. The copies must have been created after August 8, 2022:

```
bpchangeprimary -pool b_pool -sd 08/01/2022
```

In the next example, the following command promotes copy 2 of all backups of `client_a`. The copies must have been created after January 1, 2022:

```
bpchangeprimary -copy 2 -cl client_a -sd 01/01/2022
```

Duplicate backup images

NetBackup does not verify in advance whether the storage units and the drives that are required for the duplicate operation are available for use. NetBackup verifies that the destination storage units exist. The storage units must be connected to the same media server.

[Table 23-3](#) lists the scenarios in which duplication is or is not possible:

Table 23-3 Backup duplication scenarios

Duplication possible	Duplication not possible
■ From one storage unit to another. ■ From one media density to another. ■ From one server to another. ■ From multiplex to nonmultiplex format. ■ From multiplex format and retain the multiplex format on the duplicate. The duplicate can contain all or any subset of the backups that were included in the original multiplexed group. The duplicate is created with a single pass of the tape. (A multiplexed group is a set of backups that were multiplexed together during a single session.)	■ While the backup is created (unless making multiple copies concurrently). ■ When the backup has expired. ■ By using NetBackup to schedule duplications automatically (unless you use a Vault policy to schedule duplication) ■ When it is a multiplexed duplicate of the following type: ■ FlashBackup ■ NDMP backup ■ Backups from disk type storage units ■ Backups to disk type storage units ■ Nonmultiplexed backups

An alternative to duplicating backups is to create up to four copies simultaneously at backup time. (This option is sometimes referred to as Inline Copy.) Another alternative is to use storage lifecycle policies.

To duplicate backup images

- 1 On the left, click **Catalog**.
- 2 From the **Action** list, select **Duplicate**.
- 3 Select the search criteria to find the image you want to duplicate.
See [“Catalog utility search criteria and backup image details”](#) on page 395.
- 4 Select the images that you want to duplicate and click **Duplicate**.
If you duplicate a catalog backup, select all child jobs that were used to create the catalog backup. All jobs must be duplicated to duplicate the catalog backup.
- 5 Specify the number of copies you want to create. NetBackup can create up to 10 copies of unexpired backups.
If enough drives are available, the copies are created simultaneously. Otherwise, the system may require operator intervention if four copies are to be created using only two drives, for example.

- 6** The primary copy is the copy from which restores are done. Normally, the original backup is the primary copy.

If you want one of the duplicated copies to become the primary copy, select the copy number from the drop-down, otherwise select **Keep current primary copy**.

When the primary expires, a different copy automatically becomes primary. (The copy that is chosen is the one with the smallest copy number. If the primary is copy 1, copy 2 becomes primary when it expires. If the primary is copy 5, copy 1 becomes primary when it expires.)

- 7** Specify the storage unit where each copy is stored. If a storage unit has multiple drives, it can be used for both the source and destination.

All storage units must meet the criteria for creating multiple copies.

- 8** Specify the volume pool where each copy is stored.

The following volume pool selections are based on the policy type setting that was used for the query.

If the Policy type is set to All policy types (default).	Specifies that all volume pools are included in the drop-down list. Both catalog and non-catalog volume pools are included.

If the Policy type is set to NBU-Catalog .	Specifies that only catalog volume pools are included in the drop-down list.

If the Policy type is set to a policy type other than NBU-Catalog or All policy types .	Specifies that only non-catalog volume pools are included in the drop-down list.

NetBackup does not verify that the media ID selected for the duplicate copy is different from the media ID that contains the original backup. Because of this potential deadlock, specify a different volume pool to ensure that a different volume is used.

- 9** Select the retention level for the copy, or select **No change**.

The duplicate copy shares many attributes of the primary copy, including backup ID. Other attributes apply only to the primary. (For example, elapsed time.) NetBackup uses the primary copy to satisfy restore requests.

Consider the following items when selecting the retention level:

- If **No change** is selected for the retention period, the expiration date is the same for the duplicate and the source copies. You can use the `bpxupdate` command to change the expiration date of the duplicate.

- If a retention period is indicated, the expiration date for the copy is the backup date plus the retention period. For example, if a backup was created on November 14, 2022 and its retention period is one week, the new copy's expiration date is November 21, 2022.
- 10 Specify whether the remaining copies should continue or fail if the specified copy fails.
 - 11 Specify who should own the media onto which you duplicate images.
 Select one of the following:

Any	Specifies that NetBackup chooses the media owner, either a media server or server group.
None	Specifies the media server that writes to the media owns the media. No media server is specified explicitly, but you want a media server to own the media.
A server group	Specifies that only those media servers in the group are allowed to write to the media on which backup images for this policy are written. All of the media server groups that are configured in your NetBackup environment appear in the drop-down list.
 - 12 If the selection includes multiplexed backups and the backups are to remain multiplexed in the duplicate, select **Preserve multiplexing**. If you do not duplicate all the backups in a multiplexed group, the duplicate contains a different layout of fragments. (A multiplexed group is a set of backups that were multiplexed together during a single session.)

 By default, duplication is done serially and attempts to minimize media mounts and positioning time. Only one backup is processed at a time. If **Preserved multiplexing** is enabled, NetBackup first duplicates all backups that cannot be multiplex duplicated before the multiplexed backups are duplicated.

 The **Preserve multiplexing** setting does not apply when the destination is a disk storage unit. However, if the source is a tape and the destination is a disk storage unit, select **Preserve multiplexing** to ensure that the tape is read in one pass.
 - 13 Click **Yes** to start duplicating.
 - 14 Click the **Activity** tab, then select the duplication job to view the job results.
 See ["Multiplexed duplication considerations"](#) on page 403.

Multiplexed duplication considerations

Consider the following items about multiplexed duplication.

Table 23-4 Multiplexed duplication considerations

Consideration	Description
Multiplex settings are ignored	When multiplexed backups are duplicated, the multiplex settings of the destination storage unit and the original schedule are ignored. However, if multiple multiplexed groups are duplicated, the grouping within each multiplexed group is maintained. This means that the duplicated groups have a multiplexing factor that is no greater than the factor that was used during the original backup.
Backups in a multiplexed group are duplicated and duplicated group is identical	When backups in a multiplexed group are duplicated to a storage unit, the duplicated group is identical as well. However, the storage unit must have the same characteristics as the unit where the backup was originally performed. The following items are exceptions: ■ If EOM (end of media) is encountered on either the source or the destination media. ■ If any of the fragments are zero length in the source backups, the fragments are removed during duplication. A fragment of zero length occurs if many multiplexed backups start at the same time.

Jobs that appear while making multiple copies

When multiple copies are made concurrently, a parent job appears, plus a job for each copy.

The parent job displays the overall status, whereas the copy jobs display the status of a single copy. Viewing the status of individual jobs lets you troubleshoot jobs individually. For example, if one copy fails but the other copy is successful, or if each copy fails for different reasons. If at least one copy is successful, the status of the parent job is successful. Use the **Parent Job ID** filter to display the parent Job ID. Use the **Copy number** filter to display the copy number for a particular copy.

Expire backup images

To expire a backup image means to force the retention period to expire, or information about the backup is deleted. When the retention period expires,

NetBackup deletes information about the backup. The files in the backups are unavailable for restores without first re-importing.

To expire a backup image

- 1 On the left, click **Catalog**.
- 2 Select the search criteria to find the image you want to duplicate.
See [“Catalog utility search criteria and backup image details”](#) on page 395.
- 3 Select the image you want to expire and click **Expire > Expire**.

About importing backup images

NetBackup can import the backups that have expired or the backups from another NetBackup server.

During an import operation, NetBackup recreates NetBackup catalog entries for the backups on the imported volume. The import capability is useful for moving volumes from one site to another and for recreating NetBackup catalog entries.

An image is imported in the following two phases:

Table 23-5 Phases to import an image

Phase	Description
Phase I: Initiate Import	NetBackup creates a list of expired catalog entries for the backups on the imported volume. No actual import occurs in Phase I. See “Import backup images, Phase I” on page 405.
Phase II: Import	Images are selected for importing from the list of expired images that was created in Phase I. See “Import backup images, Phase II” on page 406.

About importing expired images

The expiration date for the imported items is the current date plus the retention period. For example, if a backup is imported on November 14, 2021, and its retention period is one week, the new expiration date is November 21, 2021.

Consider the following items when importing backup images:

- NetBackup can import the disk images that NetBackup version 6.0 (or later) writes.
- You cannot import a backup if an unexpired copy of it already exists on the server.

- NetBackup does not direct backups to imported volumes.
- If you import a catalog backup, import all the child jobs that were used to create the catalog backup. All jobs must be imported to import the catalog backup.
- To import a volume with the same media ID as an existing volume on a server, use the following example where you want to import a volume with media ID A00001. (A volume with media ID A00001 already exists on the server.)
 - Duplicate the existing volume on the server to another media ID (for example, B00001).
 - Remove information about media ID A00001 from the NetBackup catalog by running the following command:
 On Windows:

```
install_path\NetBackup\bin\admincmd\bpexpdate
-d 0 -m mediaID
```

 On UNIX:

```
/usr/opensv/netbackup/bin/admincmd/bpexpdate -d 0 -m
media_ID
```
 - Delete media ID A00001 from Media Manager on the server.
 - Add the other A00001 to Media Manager on the server.

To avoid this problem in the future, use unique prefix characters for media IDs on all servers.

See [“Expire backup images”](#) on page 403.

Import backup images, Phase I

Phase I of the import process creates a list of images from which to select to import in Phase II. No import occurs in Phase I.

Note the following about importing backup images:

- If tape is used, each tape must be mounted and read. It may take some time to read the catalog and build the list of images.
- The backup is not imported if it begins on a media ID that the initiating backup procedure did not process.
- The backup is incomplete if it ends on a media ID that the initiating backup procedure did not process.
- To import a catalog backup, import all of the child jobs that were used to create the catalog backup.

To perform Phase I: initialize import of backup images

- 1 To import the images from tape, make the media accessible to the media server so the images can be imported.
- 2 On the left, click **Catalog**.
- 3 On the **Actions** menu, select **Phase I import**.
- 4 For the **Media server**, specify the name of the host that contains the volume to import. This media server becomes the media owner.
- 5 Indicate the location of the image. For the **Image type**, select whether the images to be imported are located on tape or on disk.

The following table shows the actions to take depending on the location of the image.

If images are on tape	In the Media ID field, enter the Media ID of the volume that contains the backups to import.
If images are on disk	In the Disk type field, select the type of the disk storage unit on which to search for backup images. The disk types depend on which NetBackup options are licensed. If the disk type references a disk pool, enter or select the disk pool and the disk volume ID. For a BasicDisk type, enter or browse to the path to the images in the field provided. For other disk types, select <All> or the specific volume.

- 6 Click **Import** to begin reading the catalog information from the source volume.
- 7 Click on the **Activity** tab to watch as NetBackup looks at each image on the tape. NetBackup determines whether or not each image has expired and can be imported. The job also displays in the Activity monitor as an Image import type. Select the import job log to view the job results.

Import backup images, Phase II

To import the backups, first run the Initiate Import operation (Import Phase I). The first phase reads the catalog to determine all of the media that contain the catalog backup images. After Phase I, start the Import operation (Phase II). If Phase II is run before Phase I, the import fails with a message. For example, Unexpected EOF or Import of backup ID failed, fragments are not consecutive.

To import backup images, Phase II

- 1** On the left, click **Catalog**.
- 2** On the **Actions** menu, select **Phase II import**.
- 3** Set up the search criteria to find images available to import. Be sure to select a date range that includes the images you want to import. Click **Search**.
- 4** Select the images that you want to import. Click **Import** to import the selected images.
- 5** Select whether you'd like to log the names of all of the files that are found in the imported images. Click **OK**.
- 6** Click the **Activity** tab to view the progress of Import phase II.

Pausing data protection activity

This chapter includes the following topics:

- [Pause backups and other activity](#)
- [Allow the automatic pause of data protection activity](#)
- [Pause backups and other activity on a client](#)
- [View paused backups and other paused activities](#)
- [Resume data protection activity](#)

Pause backups and other activity

By default, NetBackup or its users cannot pause data protection activities. Backups and other activities continue even if a scan detects malware in an image or a recovery point. Data protection activity includes backups, duplication, and, image expiration.

You can allow NetBackup and authorized users to pause data protection activities. Then NetBackup can automatically pause activity on specific clients. For example, if a scan detects malware in backup images or recovery points for a specific client. A pause applies to scheduled backups and other automatic activities. It also applies to operations that a user initiates.

Authorized users can manually pause data protection activities. These users have an RBAC role with the necessary security permissions to pause data protection activity.

Allow the automatic pause of data protection activity

You can choose to allow NetBackup and authorized users to pause backups and duplication. Optionally, you can allow also the pause the expiration of backup images.

To allow NetBackup and authorized users to pause data protection activity

- 1 On the left, click **Detection and reporting > Paused protection**.
- 2 Click **Edit settings** and then **Edit**.
- 3 Click **Allow automatic pause**.
- 4 (Conditional) If you want to allow the pause of the expiration of backup images, select **Pause image expiration**.

Pause backups and other activity on a client

Users can pause backups and other activity on a client until a certain date or indefinitely. This functionality is available in the API endpoint `POST /config/paused-clients/`.

The following conditions occur when a client is added in the paused protection list:

- Automatic and manual replication of the client is paused.
- If the **Automatic pause protection > Pause image expiration** option is enabled, the automatic image cleanup for the client is paused.

View paused backups and other paused activities

You can view a list of the clients or hosts where data protection activity is paused.

To view paused data protection activity

- 1 On the left, click **Detection and reporting > Protection status**.
- 2 The page displays the list of clients where the protection activity is paused. "Automatic" indicates that the pause was applied automatically by NetBackup. "User-initiated" indicates that a user manually applied the pause to the client.

If you have not yet configured the setting, click **Edit settings**.
- 3 To see the details of the pause for a specific client, locate the client name. Then click **Actions > View pause details**.

Resume data protection activity

After performing maintenance or resolving any issues, you can resume the data protection activity where it is paused on a client. Perform this action from the **Detection and reporting > Paused protection** node.

Note that when you resume data protection activity, this action also turns off any host property settings that disable backups on any clients.

To resume data protection activity for a client

- 1 On the left, click **Detection and reporting > Paused protection**.
- 2 Select one or more clients and click **Resume**.

Managing security

- [Chapter 25. Security events and audit logs](#)
- [Chapter 26. Managing security certificates](#)
- [Chapter 27. Managing host mappings](#)
- [Chapter 28. Minimizing security configuration risk](#)
- [Chapter 29. Configuring multi-person authorization](#)
- [Chapter 30. Managing user sessions](#)
- [Chapter 31. Configuring multifactor authentication](#)
- [Chapter 32. Managing the global security settings for the primary server](#)
- [Chapter 33. Using access keys, API keys, and access codes](#)
- [Chapter 34. Configuring authentication options](#)
- [Chapter 35. Managing role-based access control](#)
- [Chapter 36. Disabling access to NetBackup interfaces for OS Administrators](#)

Security events and audit logs

This chapter includes the following topics:

- [View security events and audit logs](#)
- [About NetBackup auditing](#)
- [Send audit events to system logs](#)
- [Send audit events to log forwarding endpoints](#)

View security events and audit logs

NetBackup audits user-initiated actions in a NetBackup environment to help answer who changed what and when they changed it. For a full audit report, use the `nbauditreport` command. See [“Viewing the detailed NetBackup audit report”](#) on page 417.

To view security events and audit logs

- 1 On the left, select **Security > Security events**.
- 2 The following options are available.
 - Select **Access history** to view the users that accessed NetBackup.
 - Select **Audit events** to view the events that NetBackup audited. These events include changes to security settings, certificates, and users who browsed or restored backup images.

About NetBackup auditing

Auditing is enabled by default in new installations. NetBackup auditing can be configured directly on a NetBackup primary server.

Auditing of NetBackup operations provides the following benefits:

- Customers can gain insight from audit trails while they investigate unexpected changes in a NetBackup environment.
- Regulatory compliance.
The record complies with guidelines such as those required by the Sarbanes-Oxley Act (SOX).
- A method for customers to adhere to internal change management policies.
- Help for NetBackup Support in troubleshooting problems for customers.

About the NetBackup Audit Manager

The NetBackup Audit Manager (`nbaudit`) runs on the primary server and audit records are maintained in the Enterprise Media Manager (EMM) database.

An administrator can search specifically for:

- When an action occurred
- Failed actions in certain situations
- The actions that a specific user performed
- The actions that were performed in a specific content area
- Changes to the audit configuration

Note the following:

- The audit record truncates any entries greater than 4096 characters. (For example, policy name.)
- The audit record truncates any restore image IDs greater than 1024 characters.

Actions that NetBackup audits

NetBackup records the following user-initiated actions.

Activity monitor actions	Canceling, suspending, resuming, restarting, or deleting any type of job creates an audit record.
Alerts and email notifications	If an alert cannot be generated or an email notification cannot be sent for NetBackup configuration settings. For example, SMTP server configuration and the list of excluded status codes for alerts.

Anomalies	When a user reports an anomaly as false positive, the action is audited and logged for that user.
Malware detection	When malware scan is triggered, malware scan status and malware scan configuration actions are audited.
Asset actions	Deleting an asset, such as a vCenter server, as part of the asset cleanup process is audited and logged. Creating, modifying, or deleting an asset group as well any action on an asset group for which a user is not authorized is audited and logged.
Authorization failure	Authorization failure is audited when you use the NetBackup web UI, or the NetBackup APIs.
Catalog information	This information includes: ■ Verifying and expiring images. ■ Read the requests that are sent for the front-end usage data.
Certificate management	Creating, revoking, renewing, and deploying of NetBackup certificates and specific NetBackup certificate failures.
Certificate Verification Failures (CVFs)	Any failed connection attempts that involve SSL handshake errors, revoked certificates, or host name validation failures. For certificate verification failures (CVFs) that involve SSL handshakes and revoked certificates, the timestamp indicates when the audit record is posted to the primary server. (Rather than when an individual certificate verification fails.) A CVF audit record represents a group of CVF events over a time period. The record details provide the start and the end times of the time period as well as the total number of CVFs that occurred in that period.
Disk pools and Volume pools actions	Adding, deleting, or updating disk or volume pools.
Hold operations	Creating, modifying, and deleting hold operations.
Host database	NetBackup operations that are related to the host database.
IRE configuration and states	Adding, updating, and deleting IRE allowed subnets or schedule. IRE external network is opened or closed by IRE schedule or by an administrator.
Logon attempts	Any successful or any failed logon attempts for the NetBackup web UI or the NetBackup APIs.
Policies actions	Adding, deleting, or updating policy attributes, clients, schedules, and backup selections lists.

Restore and browse image user actions	All the restore and browse image content (<code>bplist</code>) operations that a user performs are audited with the user identity. To set an interval to periodically add audit records of the browse image (<code>bplist</code>) operations from the cache into the NetBackup database, use the <code>DATAACCESS_AUDIT_INTERVAL_HOURS</code> configuration option. Setting this configuration option prevents the NetBackup database size from increasing exponentially because of the <code>bplist</code> audit records. See the NetBackup Administrator's Guide Volume I. To add all the <code>bplist</code> audit records from the cache into the NetBackup database, run the following command on the primary server: <pre>nbcertcmd -postAudit -dataAccess</pre>
Security configuration	Information that is related to changes that are made to the security configuration settings.
Starting a restore job	NetBackup does not audit when other types of jobs begin. For example, NetBackup does not audit when a backup job begins.
Starting and stopping the NetBackup Audit Manager (<code>nbaudit</code>).	Starting and stopping of the <code>nbaudit</code> manager is always audited, even if auditing is disabled.
Storage lifecycle policy actions	Attempts to create, modify, or delete a storage lifecycle policy (SLP) are audited and logged. However, activating and suspending an SLP using the command <code>nbstlutil</code> are not audited. These operations are audited only when they are initiated from a NetBackup graphical user interface or API.
Storage servers actions	Adding, deleting, or updating storage servers.
Storage units actions	Adding, deleting, or updating storage units. Note: Actions that are related to storage lifecycle policies are not audited.
Token management	Creating, deleting, and cleanup of tokens and specific token issuing failures.
User action that fails to create an audit record	If auditing is enabled but a user action fails to create an audit record, the audit failure is captured in the <code>nbaudit</code> log. NetBackup status code 108 is returned (<code>Action succeeded but auditing failed</code>). The NetBackup does not return an exit status code 108 when auditing fails.

Actions that NetBackup does not audit

The following actions are not audited and do not display in the audit report:

Any failed actions.	NetBackup logs failed actions in NetBackup error logs. Failed actions do not display in audit reports because a failed attempt does not bring about a change in the NetBackup system state.

The effect of a configuration change	The results of a change to the NetBackup configuration are not audited. For example, the creation of a policy is audited, but the jobs that result from its creation are not.
The completion status of a manually initiated restore job	While the act of initiating a restore job is audited, the completion status of the job is not audited. Nor is the completion status of any other job type, whether initiated manually or not. The completion status is displayed in the Activity Monitor.
Internally initiated actions	NetBackup-initiated internal actions are not audited. For example, the scheduled deletion of expired images, scheduled backups, or periodic image database cleanup is not audited.
Rollback operations	Some operations are carried out as multiple steps. For example, creating an MSDP-based storage server consists of multiple steps. Every successful step is audited. Failure in any of the steps results in a rollback, or rather, the successful steps may need to be undone. The audit record does not contain details about rollback operations.
Host properties actions	Changes made with the <code>bpsetconfig</code> or the <code>nbsetconfig</code> commands, or the equivalent property in host properties, are not audited. Changes that are made directly to the <code>bp.conf</code> file or to the registry are not audited.

User identity in the audit report

The audit report indicates the identity of the user who performed a specific action. The full identity of the user includes the user name and the domain or the host name that is associated with the authenticated user. A user's identity appears in the audit report as follows:

- Audit events always include the full user identity. Root users and administrators are logged as "root@hostname" or "administrator@hostname".
- In NetBackup 8.1.2 and later, image browse and image restore events always include the user ID in the audit event. NetBackup 8.1.1 and earlier log these events as "root@hostname" or "administrator@hostname".
- The order of the elements for the user principal is "domain:username:domainType:providerId". The domain value does not apply for Linux computers. For that platform, the user principal is :username:domainType:providerId.
- For any operations that do not require credentials or require the user to sign in, operations are logged without a user identity.

Audit retention period and catalog backups of audit records

The audit records are kept as part of the NetBackup database, for as long as the retention period indicates. The records are backed up as part of the NetBackup catalog backup. The NetBackup Audit Service (`nbaudit`) deletes expired audit records once every 24 hours at 12:00 A.M. (local time).

If no audit retention period is indicated, audit records are retained for 90 days, which is a default value. Set the audit retention period to 0 (zero) if you do not want to delete the audit records.

To configure the audit retention period

1 Log on to the primary server.

2 Run the following command:

```
bpnbat -login
```

3 Open the following directory:

Windows: `install_path\NetBackup\bin\admincmd`

UNIX: `/usr/opensv/netbackup/bin/admincmd`

4 Enter the following command:

```
nbseccmd -setsecurityconfig -auditretentionperiod number_of_days
```

The audit report is retained for the value that is specified for the *number_of_days* option.

In the following example, the records of user actions are retained for 30 days and then deleted.

```
nbseccmd -setsecurityconfig -auditretentionperiod 30
```

To ensure that audit records are backed up during catalog backups, configure the catalog backup frequency to be less frequent or equal to the value that you specify for `-auditretentionperiod`.

5 To check the current audit retention period, run the following command:

```
nbseccmd -getsecurityconfig -auditretentionperiod
```

Viewing the detailed NetBackup audit report

You can view the actions NetBackup audits from a primary server using the NetBackup web UI. You can see full audit event details with the `nbauditreport` command.

To view the full audit report

- 1 Log on to the primary server.
- 2 Enter the following command to display the audit report in the summary format.

Windows: `install_path\NetBackup\bin\admincmd\nbauditreport`

UNIX: `/usr/opensv/netbackup/bin/admincmd\nbauditreport`

Or, run the command with the following options.

`-sdate`
`<"MM/DD/YY`
`[HH:[MM[:SS]]]">`

The start date and time of the report data you want to view.

`-edate`
`<"MM/DD/YY`
`[HH:[MM[:SS]]]">`

The end date and time of the report data you want to view.

`-ctgy category`

The category of user action that was performed. Categories such as `POLICY` may contain several sub-categories such as schedules or backup selections. Any modifications to a sub-category are listed as a modification to the primary category.

See the [NetBackup Commands Guide](#) for `-ctgy` options.

`-user`
`<username[:domainname]>`

Use to indicate the name of the user for whom you'd like to display audit information.

`-fmt DETAIL`

The `-fmt DETAIL` option displays a comprehensive list of audit information. For example, when a policy is changed, this view lists the name of the attribute, the old value, and the new value. This option has the following sub-options:

- `[-notruncate]` . Display the old and new values of a changed attribute on separate lines in the details section of the report.
- `[-pagewidth <NNN>]` . Set the page width for the details section of the report.

`-fmt PARSABLE`

The `-fmt PARSABLE` option displays the same set of information as the `DETAIL` report but in a parsable format. The report uses the pipe character (`|`) as the parsing token between the audit report data. This option has the following sub-options:

- `[-order <DTU|DUT|TDU|TUD|UDT|UTD>]`.
Indicate the order in which the information appears.
 - D (Description)
 - T (Timestamp)
 - U (User)

3 The audit report contains the following details:

DESCRIPTION	The details of the action that was performed.
USER	The identity of the user who performed the action. See "User identity in the audit report" on page 416.
TIMESTAMP	The time that the action was performed.

The following information only displays if you use the `-fmt DETAIL` or the `-fmt PARSABLE` options.

CATEGORY	The category of user action that was performed.
ACTION	The action that was performed.
REASON	The reason that the action was performed. A reason displays if a reason was specified for the operation that created the change.
DETAILS	An account of all of the changes, listing the old values and the new values.

Example of the audit report:

```
[root@server1 admincmd]# ./nbauditreport
TIMESTAMP          USER              DESCRIPTION
04/20/2018 11:52:43 root@server1      Policy 'test_pol_1' was saved but no changes were detected
04/20/2018 11:52:42 root@server1      Schedule 'full' was added to Policy 'test_pol_1'
04/20/2018 11:52:41 root@server1      Policy 'test_pol_1' was saved but no changes were detected
04/20/2018 11:52:08 root@server1      Policy 'test_pol_1' was created
04/20/2018 11:17:00 root@server1      Audit setting(s) of master server 'server1' were modified
```

Audit records fetched: 5

Send audit events to system logs

You can send NetBackup audit events to system logs. You must have the NetBackup Security Administrator role or similar RBAC permissions to perform this task.

By default, NetBackup sends the audit events to system logs in native format. You can now export audit events with the Open Cybersecurity Schema Framework (OCSF) format to Security Information and Event Management (SIEM) platforms.

See [this article](#) for more information.

Use the `SYSLOG_AUDIT_USE_OCSF_FORMAT` configuration option to send the NetBackup audit events to system logs in the OCSF format.

To send audit events to system logs

- 1 Open the NetBackup web UI.
- 2 On the left, select **Security > Security events**.
- 3 On the top right, click **Security event settings**.
- 4 Enable the **Send the audit events to the system logs** option.
- 5 Select **Select audit event** categories. Then select the audit categories for which you want to send the audit events to the system logs.

To send audit events for all audit categories to the system logs, select the **Audit event categories** check box.

- 6 Select **Save**.

You can view NetBackup audit events in the system logs. For example:

On a Windows system, use **Windows Event Viewer** to view NetBackup audit events.

On a Linux system, you can view the system logs on the configured location.

Send audit events to log forwarding endpoints

You can send NetBackup audit events to log forwarding endpoints.

To send audit events to log forwarding endpoints

- 1 On the left, select **Security > Security events**.
- 2 On the top right, select **Security events settings**.
- 3 Enable the **Send the audit events to log forwarding endpoints** option.
After you enable the option, the **Select endpoints and categories** option displays.
- 4 Select the **Select endpoints and categories** option to see the log forwarding endpoints that are configured in your environment and the available audit categories.
Example of an endpoint: Azure Sentinel.
- 5 Select the appropriate log forwarding endpoints.
- 6 Select the **Select audit event categories** option.
- 7 Select the categories of the audit events that you want to forward to the selected endpoints. For example, Alert, Anomaly, etc.
- 8 After you select your log forwarding endpoint, the options to specify the associated credentials display. You can either add new credentials for the endpoint or select the existing credentials.

Managing security certificates

This chapter includes the following topics:

- [About security management and certificates in NetBackup](#)
- [NetBackup host IDs and host ID-based certificates](#)
- [Manage NetBackup security certificates](#)
- [Using external security certificates with NetBackup](#)

About security management and certificates in NetBackup

NetBackup uses security certificates to authenticate the NetBackup hosts. These certificates must conform to the X.509 public key infrastructure (PKI) standard. With NetBackup 8.1, 8.1.1, and 8.1.2, NetBackup certificates are used for secure communication. In NetBackup 8.2 and later you can use NetBackup certificates or external certificates.

NetBackup certificates are issued to hosts by default and the NetBackup primary server acts as the CA and manages the Certificate Revocation List (CRL). The **NetBackup certificate deployment security level** determines how certificates are deployed to NetBackup hosts and how often the CRL is updated on each host. If a host needs a new certificate (the original certificate is expired or revoked), you can use an NetBackup authorization token to reissue the certificate.

External certificates are those that a trusted external CA signed. When you configure NetBackup to use external certificates, the primary server, media servers, and clients in the NetBackup domain use the external certificates for secure

communication. Additionally, the NetBackup web server uses these certificates for communication between the NetBackup web UI and the NetBackup hosts. Deployment of external certificates, updating or replacing external certificates, and CRL management for the external CA are managed outside of NetBackup.

For more information on external certificates, see the [NetBackup Security and Encryption Guide](#).

Security certificates for NetBackup 8.1 and later hosts

NetBackup 8.1 and later hosts can communicate with each other only in a secure mode. Depending on the NetBackup version, these hosts must have a certificate that the NetBackup CA issued or that another trusted CA issued. A NetBackup certificate that is used for secure communications over a control channel is also referred to as host ID-based certificate.

Security certificates for NetBackup 8.0 hosts

Any security certificates that NetBackup generated for 8.0 hosts are referred to as host name-based certificates. For more details on these certificates, refer to the [NetBackup Security and Encryption Guide](#).

NetBackup host IDs and host ID-based certificates

Each host in a NetBackup domain has a unique identity, which is referred to as a host ID or a Universally Unique Identifier (UUID). The host ID is used in many operations to identify the host. NetBackup creates and manages host IDs as follows:

- Maintains a list on the primary server of all of the host IDs that have certificates.
- Randomly generates host IDs. These IDs are not tied to any property of the hardware.
- By default, assigns NetBackup 8.1 and later hosts a host ID-based certificate that is signed by the NetBackup certificate authority.
- The host ID remains the same even when the host name changes.

In some cases a host can have multiple host IDs:

- If a host obtains certificates from multiple NetBackup domains, it has multiple host IDs that correspond to each NetBackup domain.
- When the primary server is configured as part of a cluster, each node in the cluster receives a unique host ID. An additional host ID is assigned for the virtual name. For example, if the primary server cluster is composed of N nodes, the number of host IDs that are allocated for the primary server cluster is $N + 1$.

Manage NetBackup security certificates

Note: The information here only applies to the security certificates that the NetBackup certificate authority (CA) issues. More information is available for external certificates.

See [“Using external security certificates with NetBackup”](#) on page 428.

You can view and revoke NetBackup certificates and view information about the NetBackup CA. More detailed information about NetBackup certificate management and certificate deployment is available in the [NetBackup Security and Encryption Guide](#).

View a NetBackup certificate

You can view details of all host ID-based NetBackup certificates that are issued to NetBackup hosts. Note that only 8.1 and later NetBackup hosts have host ID-based certificates. The **Certificates** list does not include any NetBackup 8.0 or earlier hosts.

To view a NetBackup certificate

- 1 On the left, select **Security > Certificates**.
- 2 Select the **NetBackup certificates** tab.
- 3 To view additional certificate details for a host, click on a host name.

Revoke a NetBackup CA certificate

When you revoke a NetBackup host ID-based certificate, NetBackup revokes any other certificates for that host. NetBackup ceases to trust the host, and it can no longer communicate with the other NetBackup hosts.

You can choose to revoke a host ID-based certificate under various conditions. For example, if you detect that client security has been compromised, if a client is decommissioned, or if NetBackup was uninstalled from the host. A revoked certificate cannot be used to communicate with primary server web services.

Security best practices suggest that the NetBackup security administrator explicitly revoke the certificates for any host that is no longer active. Take this action if whether or not the certificate is still deployed on the host.

Note: Do not revoke a certificate of the primary server. If you do, NetBackup operations may fail.

To revoke a NetBackup CA certificate

- 1 On the left, select **Security > Certificates**.
- 2 Select the **NetBackup certificates** tab.
- 3 Select the host that is associated with the certificate that you want to revoke.
- 4 Select **Revoke certificate > Yes**.

View the NetBackup certificate authority details and fingerprint

For secure communication with the NetBackup certificate authority (CA) on the primary server, a host's administrator must add the CA certificate to an individual host's trust store. The primary server administrator must give the fingerprint of the CA certificate to the administrator of the individual host.

To view the NetBackup certificate authority details and fingerprint

- 1 On the left, select **Security > Certificates**.
- 2 Click the **NetBackup certificates** tab.
- 3 In the toolbar, select **Certificate authority**.
- 4 Find the **Fingerprint** information and select **Copy to clipboard**.
- 5 Provide this fingerprint information to the host's administrator.

Reissue a NetBackup certificate

Note: The information here only applies to the security certificates that the NetBackup certificate authority (CA) issues. External certificates must be managed outside of NetBackup.

In some cases a host's NetBackup certificate is no longer valid. For example, if a certificate is expired, revoked, or is lost. You can reissue a certificate either with or without a reissue token.

A reissue token is a type of authorization token that is used to reissue a NetBackup certificate. When you reissue a certificate, the host gets the host ID same as the original certificate.

Reissue a NetBackup certificate, with a token

If you need to reissue a host's NetBackup certificate NetBackup provides a more secure method to do this reissue. You can create an authorization token that the

host administrator must use to obtain a new certificate. This reissue token retains the same host ID as the original certificate. The token can only be used once. Because it is associated to a specific host, the token cannot be used to request certificates for other hosts.

To reissue a NetBackup certificate for a host

- 1 On the left, select **Security > Certificates**.
- 2 Select the **NetBackup certificates** tab.
- 3 Select the host and select **Actions > Generate reissue token**.
- 4 Enter a token name and indicate how long the token should be valid for.
- 5 Select **Create**.
- 6 Select **Copy to clipboard** and then select **Close**.
- 7 Share the authorization token so the host's administrator can obtain a new certificate.

Allow a NetBackup certificate reissue, without a token

In certain scenarios you need to reissue a certificate without a reissue token. For example, for a BMR client restore. The option **Allow auto reissue certificate** enables you to reissue a certificate without requiring a token.

To allow a NetBackup certificate reissue, without a token

- 1 On the left, select **Security > Host mappings**.
- 2 Locate the host and select **Actions > Allow auto reissue certificate > Allow**.

Once you set the **Allow auto reissue certificate** option, a certificate can be reissued without a token within the next 48 hours, which is the default setting. After this window to reissue expires, the certificate reissue operation requires a reissue token.
- 3 Notify the host's administrator that you allowed a NetBackup certificate reissue without a token.

Revoke the ability to reissue a NetBackup certificate without a token

After you allow a NetBackup certificate reissue without a token, you can revoke this ability before the window to reissue expires. By default, the window is 48 hours.

To revoke the ability to reissue a NetBackup certificate without a token

- 1 On the left, select **Hosts > Host mappings**.
- 2 Locate the host and select **Actions > Revoke auto reissue certificate > Revoke**.

Manage NetBackup certificate authorization tokens

Note: The information here only applies to the security certificates that the NetBackup certificate authority (CA) issues. External certificates must be managed outside of NetBackup.

Depending on the security level for NetBackup certificate deployment, you may need an authorization token to issue a new NetBackup certificate to a host. You can create a token when it is required or find and copy a token if it is needed again. Tokens can be cleaned up or deleted if they are no longer needed.

To reissue a certificate, a reissue token is required in most cases. A reissue token is associated with the host ID.

Create an authorization token

Depending on the NetBackup certificate deployment security level, an authorization token may be required for a non-primary NetBackup host to obtain a host ID-based NetBackup certificate. The NetBackup administrator of the primary server generates the token and shares it with the administrator of the non-primary host. That administrator can then deploy the certificate without the presence of the primary server administrator.

Do not create an authorization token for a NetBackup host whose current certificate is not in a valid state because it is lost, corrupt, or expired. In these cases, a reissue token must be used.

See [“Reissue a NetBackup certificate”](#) on page 425.

To create an authorization token

- 1 On the left, select **Security > Tokens**.
- 2 On the top left, select **Add**.
- 3 Enter the following information for the token:
 - Token name
 - The maximum number of times you want the token to be used
 - How long the token is valid for
- 4 Select **Create**.

To find and copy an authorization token value

You can view the details of the tokens that you have created and copy the token value for future use.

To find and copy an authorization token value

- 1 On the left, select **Security > Tokens**.
- 2 Select the name of the token for which you want to view the details.
- 3 Select **Show** and then click the **Copy to clipboard** icon.

Cleanup tokens

Use the Cleanup tokens utility to delete tokens from the token database that are expired or that have reached the maximum number of uses allowed.

To cleanup tokens

- 1 On the left, select **Security > Tokens**.
- 2 Click **Cleanup > Yes**.

Delete a token

You can delete a token before it is expired or before the **Maximum uses allowed** is reached.

To delete a token

- 1 On the left, select **Security > Tokens**.
- 2 Select the name of the tokens that you want to delete.
- 3 Select **Delete**.

Using external security certificates with NetBackup

NetBackup 8.2 and later versions support the security certificates that are issued by an external CA. External certificates and the certificate revocation list for an external certificate authority must be managed outside of NetBackup. The **External certificates** tab displays details for the NetBackup 8.1 and later hosts in the domain and whether or not they use external certificates.

Before you can see external certificate information in **Certificates > External certificates**, you must first configure the primary server and the NetBackup web server to use external certificates.

See [“Configure an external certificate for the NetBackup web server”](#) on page 429.

See the video [External CA support in NetBackup](#).

Configure an external certificate for the NetBackup web server

By default, NetBackup uses the security certificates that the NetBackup CA has issued. If you have a certificate that an external CA has issued, you can configure the NetBackup web server to use it for secure communication.

Note: Windows certificate store is not supported as certificate source for the NetBackup web server.

The API that you can use to configure the external certificate for the NetBackup web server: `POST security/web-certificates/{certificate_id}`.

If external certificate for the web server is configured using the API, the configuration process is audited.

To configure an external certificate for the web server

- 1 Ensure that you have valid certificate, private key of the certificate, and trusted CA bundle.
- 2 Ensure that the NetBackup Web Management Console service is up and running.
- 3 Run the following command:

```
configureWebServerCerts -addExternalCert -nbHost -certPath
certificate path -privateKeyPath private key path -trustStorePath
CA bundle path [-passphrasePath passphrase file path]
```

The `configureWebServerCerts` command does not support use of Windows certificate store paths.

Refer to the [NetBackup Commands Reference Guide](#) for more details on the command-line options.

- In a clustered setup, to avoid a failover run the following command on the active node:

```
install_path/netbackup/bin/bpclusterutil -freeze
```

- If the FIPS mode is enabled on the primary server, you can use only the PEM-formatted files for the `configureWebServerCerts` command.

- 4 Restart the NetBackup Web Management Console service to reflect the changes.

On UNIX, run the following commands:

- `install_path/netbackup/bin/nbwmc -terminate`
- `install_path/netbackup/bin/nbwmc start`

On Windows, use the **Services** application in the **Windows Control Panel**.

Location of the commands:

Windows `install_path\NetBackup\wmc\bin\install\`

UNIX `install_path/wmc/bin/install`

- In a clustered setup, unfreeze the cluster using the following command on the active node:

```
install_path/netbackup/bin/bpclusterutil -unfreeze
```

- 5 Restart the NetBackup Messaging Queue Broker (`nbmqbroker`) service as follows:

On Windows:

Go to the **Services** application in the **Windows Control Panel** and manually restart the NetBackup Messaging Queue Broker service.

On UNIX:

Run the following command:

```
nbmqbroker stop; nbmqbroker start
```

- 6 Verify that you can access the NetBackup web user interface using a browser, without a certificate warning message.

Remove the external certificate configured for the web server

You can remove the external certificate that is configured for the NetBackup web server. NetBackup then uses the NetBackup CA-signed certificate for secure communication.

The API that you can use to remove the external certificate for the NetBackup web server: `DELETE security/web-certificates/{certificate_id}`.

To remove the external certificate configured for the web server

- 1 Ensure that the NetBackup Web Management Console service is up and running.
- 2 Run the following command (in a clustered primary server setup, run this command on the active node):

```
configureWebServerCerts -removeExternalCert -nbHost
```

Refer to the [NetBackup Commands Reference Guide](#) for more details on the command-line options.

- In a clustered primary server setup, run the following command on the active node to freeze the cluster to avoid a failover:

```
install_path/netbackup/bin/bpclusterutil -freeze
```

3 Restart the NetBackup Web Management Console service.

- In a clustered primary server setup, run the following command on the active node to unfreeze the cluster:

```
install_path/netbackup/bin/bpclusterutil -unfreeze
```

4 Restart the NetBackup Messaging Queue Broker (`nbmqbroker`) service as follows:

On Windows:

Go to the **Services** application in the **Windows Control Panel** and manually restart the NetBackup Messaging Queue Broker service.

On UNIX:

Run the following command:

```
nbmqbroker stop; nbmqbroker start
```

Update or renew the external certificate for the web server

You can update or renew the external certificate that you configured for the web server.

To update or renew the external certificate for the web server

- 1** Ensure that you have the latest external certificate, the matching private key, and the CA bundle file.
- 2** Run the following command (in a clustered setup, run the command on the active node):

```
configureWebServerCerts -addExternalCert -nbHost -certPath  
certificate path -privateKeyPath private key path -trustStorePath  
CA bundle path
```

View external certificate information for the NetBackup hosts in the domain

Note: Before you can see external certificate information, you must configure NetBackup for external certificates. See the [NetBackup Security and Encryption Guide](#) for details.

As you add external certificates to the hosts in the NetBackup domain, use the **External certificates** dashboard to track which hosts need attention. To support an external certificate, a host must be upgraded and enrolled with an external certificate.

To view external certificate information for the hosts

- 1 On the left, select **Security > Certificates**.
- 2 Select the **External certificates** tab.

In addition to hosts information and details for the hosts' external certificates, the following information is also included:

- The **NetBackup certificate status** column indicates if a host also has a NetBackup certificate.
- The **External certificate** dashboard contains the following information for NetBackup 8.1 and later hosts:
 - Total hosts. The total number of hosts. The hosts must be online and able to communicate with NetBackup primary server.
 - Hosts with certificate. The number of hosts that have a valid external certificate enrolled with the NetBackup primary server.
 - Hosts with no certificate. Either the host supports external certificates, but does not have one enrolled. Or, an upgrade to NetBackup 8.2 or later is required for the host (applies to versions 8.1, 8.1.1, or 8.1.2). The **NetBackup upgrade required** total also includes any hosts that were reset or any hosts for which the NetBackup version is unknown. NetBackup 8.0 and earlier hosts do not use security certificates and are not reflected here.
 - Certificate expiry. The hosts that have an expired or expiring external certificate.

View details for a host's external certificate

You can view details of a host's certificate that was issued by an external certificate authority.

To view details for a host's external certificate

- 1 On the left, select **Security > Certificates**.
- 2 Click the **External certificates** tab.

The list of external certificates displays for the primary server.
- 3 To view additional certificate details for a host, click on a host name.

Managing host mappings

This chapter includes the following topics:

- [View host security and mapping information](#)
- [Approve or add mappings for a host that has multiple host names](#)
- [Example host mappings](#)
- [Remove mappings for a host that has multiple host names](#)

View host security and mapping information

The **Hosts** information in **Host mappings** contains details about the NetBackup hosts in your environment, including the primary server, media servers, and clients. Only hosts with a host ID are displayed in this list. The **Host** name reflects the NetBackup client name of a host, also referred to as the primary name of the host.

Note: NetBackup discovers any dynamic IP addresses (DHCP or Dynamic Host Configuration Protocol hosts) and adds these addresses to a host ID. You should delete these mappings.

For host name-based certificates for 8.0 and earlier NetBackup hosts, refer to the respective version of the [NetBackup Security and Encryption Guide](#).

To view NetBackup host information

- 1 On the left, select **Security > Host mappings**.
Review the security status and any other host names that are mapped to this host.
- 2 For additional details for this host, click the name of the host.

Approve or add mappings for a host that has multiple host names

A NetBackup host can have multiple host names. For example, both a private and a public name or a short name and a fully qualified domain name (FQDN). A NetBackup host may also share a name with other NetBackup host in the environment. NetBackup also discovers cluster names, including the host name and fully qualified domain name (FQDN) of the virtual name of the cluster.

The NetBackup client name of a host (or the primary name) is automatically mapped to its host ID during certificate deployment. For successful communication between NetBackup hosts, NetBackup also automatically maps all hosts to their other host names. However, that method is less secure. Instead, you can choose to disable this setting. Then choose to manually approve the individual host name mappings that NetBackup discovers.

See [“Disable automatic mapping of NetBackup host names”](#) on page 474.

See [“Example host mappings”](#) on page 436.

Approve the host mappings that NetBackup discovers

NetBackup automatically discovers many shared names or cluster names that are associated with the NetBackup hosts in your environment. Use the **Mappings to approve** tab to review and accept the relevant host names. When option **Automatically map host names to their NetBackup host ID** is enabled, the **Mappings to approve** list shows only the mappings that conflict with other hosts.

Note: You must map all available host names with the associated host ID. When you deploy a certificate to a host, the host name must map to the associated host ID. If it does not, NetBackup considers the host to be a different host. NetBackup then deploys a new certificate to the host and issues it a new host ID.

To approve the host names that NetBackup discovers

- 1 On the left, select **Security > Host mappings**.
- 2 Select the **Mappings to approve** tab.
- 3 Select the name of the host.
- 4 Review the mappings for the host and select the **Approve** button if you want to use the discovered mapping.

Select **Reject** if you do not want to associate the mapping with the host.

The rejected mappings do not appear in the list until NetBackup discovers them again.

- 5 Select the **Save** button.

Map other host names to a host

You can manually map the NetBackup host to its host names. This mapping ensures that NetBackup can successfully communicate with the host using the other name.

To map a host name to a host

- 1 On the left, select **Security > Host mappings**.
- 2 Select the host and select the **Manage mappings** button.
- 3 Select the **Add** button.
- 4 Enter the host name or IP address and select **Save**.
- 5 Select **Close**.

Map shared or cluster names to multiple NetBackup hosts

Add a shared or a cluster name mapping if multiple NetBackup hosts share a host name. For example, a cluster name.

Note the following before you create a shared or a cluster name mapping:

- NetBackup automatically discovers many shared names or cluster names. Review the **Mappings to approve** tab.
- If a mapping is shared between an insecure and a secure host, NetBackup assumes that the mapping name is secure. However, if at run-time the mapping resolves to an insecure host, the connection fails. For example, assume that you have a two-node cluster with a secure host (node 1) and an insecure host (node 2). In this case, the connection fails if node 2 is the active node.

To map shared or cluster names to multiple NetBackup hosts

- 1 On the left, select **Security > Host mappings**.
- 2 Select the **Add shared or cluster mappings** button.
- 3 Enter a **Shared host name or cluster name** that you want to map to two or more NetBackup hosts.

For example, enter a cluster name that is associated with NetBackup hosts in your environment.

- 4 On the right, select the **Add** button.

- 5 Select the NetBackup hosts that you want to add and select **Add to list**.
For example, if you entered a cluster name in step 3 select the nodes in the cluster here.
- 6 Select **Save**.

Example host mappings

The following examples describe scenarios where you may want to create host mappings to consolidate host names or to ensure successful communication between hosts.

- See [the section called “Examples of auto-discovered mappings for a cluster”](#) on page 436.
- See [the section called “Example of host names that are displayed for a multiple NIC environment”](#) on page 437.
- See [the section called “Example of auto-discovered mappings for a cluster in a multiple NIC environment”](#) on page 438.
- See [the section called “Examples of auto-discovered mappings for SQL Server environments”](#) on page 438.

Examples of auto-discovered mappings for a cluster

For a cluster with hosts `client01.lab04.com` and `client02.lab04.com`, you may see the following entries. For each host, approve the mappings that are valid.

Host	Auto-discovered mapping
client01.lab04.com	client01
client01.lab04.com	clustername
client01.lab04.com	clustername.lab04.com
client02.lab04.com	client02
client02.lab04.com	clustername
client02.lab04.com	clustername.lab04.com

After you approve all the valid mappings, you see the **Mapped host or IP address** settings that are similar to the following entries.

Host	Mapped Host Names/IP Addresses
client01.lab04.com	client01.lab04.com, client01, clustername, clustername.lab04.com
client02.lab04.com	client02.lab04.com, client02, clustername, clustername.lab04.com

Example of host names that are displayed for a multiple NIC environment

In some advanced NetBackup configurations like a multi-NIC environment, a NetBackup host may display under two host names in the **Host properties**. One name reflects the operating system (OS) name and the other name reflects the name that was specified when NetBackup was installed. This behavior does not affect the ability to connect to the host or to view or edit the host's properties.

For example, you may see the following entries for *Host 1* that is in a multi-NIC environment.

Table 27-1 Multiple host name entries for a host in a multi-NIC environment

Host	Mapped host names
osname-host1.domain.com	OS name of <i>Host 1</i>
clientname-host1.domain.com	Client name of <i>Host 1</i>

To consolidate these host names, to the host `clientname-host1.domain.com` add a mapping for `osname-host1.domain.com`. After you add the mapping, you see only one entry for the host in host properties.

Table 27-2 Host mapping for a multi-NIC environment

Host	Mapped host names
client01-name.domain.com	clientname-host1.domain.com, osname-host1.domain.com

Example of auto-discovered mappings for a cluster in a multiple NIC environment

Backups of a cluster in a multi-NIC environment require special mappings. You must map the cluster node names to the virtual name of the cluster on the private network.

Table 27-3 Mapping host names for a cluster in a multi-NIC environment

Host	Mapped host names
Private name of <i>Node 1</i>	Virtual name of the cluster on the private network
Private name of <i>Node 2</i>	Virtual name of the cluster on the private network

For example, for a cluster in a multi-NIC environment with hosts `client01-bk.lab04.com` and `client02-bk.lab04.com`, you may see the following entries. For each host, approve the mappings that are valid.

Host	Auto-discovered mapping
client01-bk.lab04.com	clustername-bk.lab04.com
client02-bk.lab04.com	clustername-bk.lab04.com

After you approve all the valid mappings, you see the **Mapped host or IP address** settings that are similar to the following entries.

Host	Mapped host names or IP addresses
client01-bk.lab04.com	clustername-bk.lab04.com
client02-bk.lab04.com	clustername-bk.lab04.com

Examples of auto-discovered mappings for SQL Server environments

In [Table 27-4](#), FCI is a SQL Server failover cluster instance. WSFC is Windows Server Failover Cluster.

Table 27-4 Example mapped host names for SQL Server environments

Environment	Host	Mapped host names
FCI (cluster with two nodes)	Physical name of <i>Node 1</i>	Virtual name of the SQL Server cluster
	Physical name of <i>Node 2</i>	Virtual name of the SQL Server cluster
Basic or advanced availability group (primary and secondary)	Primary name	WSFC name
	Secondary name	WSFC name
Basic or advanced availability group, with an FCI (primary FCI and secondary FCI)	Primary FCI name	WSFC name
	Secondary FCI name	WSFC name
	Physical name of <i>Node 1</i>	Virtual name of the SQL Server cluster
	Physical name of <i>Node 2</i>	Virtual name of the SQL Server cluster

Remove mappings for a host that has multiple host names

You can remove any host name mappings that NetBackup added automatically. Or, any host name mappings that you added manually for a host. Note that if you remove a mapping, the host is no longer recognized with that mapped name. If you remove a shared or a cluster mapping, the host may not be able to communicate with other hosts that use that shared or cluster name.

If you have issues with a host and its mappings, you can reset the host attributes. However, that resets other attributes like a host's communication status.

See [“Reset a host's attributes”](#) on page 96.

To remove a host name that NetBackup discovers

- 1 On the left, select **Security > Host mappings**.
- 2 Locate the host that you want to update.

- 3** Click **Actions > Manage mappings**.
- 4** Locate the mapping you want to remove and click **Delete > Save**.

Minimizing security configuration risk

This chapter includes the following topics:

- [About security configuration risk](#)
- [Security settings to be configured to minimize risk](#)
- [Set the current posture as security baseline](#)
- [Manage security baseline](#)
- [Manage security baseline from Alta View UI](#)

About security configuration risk

Security configuration risk depends on the status of the security settings in your NetBackup domain. A high configuration risk score implies that more number of security settings need to be configured in the domain. To minimize the risk, enable all the required security settings.

See [“Security settings to be configured to minimize risk”](#) on page 443.

The security risk score is also determined based on the active status of each host in the domain. A host is considered to be active if it has participated in secure communication within the domain during the last seven days.

Risk score for the following settings is determined based on the active status of the hosts:

- Secure data-in-transit encryption (DTE)
- Service user configuration
- Percent of servers with the current primary server version or later

- Percent of other hosts with current primary server version or later

For more information on how to minimize the security configuration risk, see the [article](#).

The NetBackup web UI dashboard shows the security configuration risk score. When you change any of the security settings, the risk score is updated on the dashboard.

The following parameters help you learn the current security scenario in your domain and how you can minimize the security configuration risk.

Use the NetBackup web UI dashboard to see these parameters.

Current posture

Current posture comprises the current values of NetBackup security settings. It is recommended that you enable all security settings to minimize the security configuration risk.

See [“Security settings to be configured to minimize risk”](#) on page 443.

Security baseline

Security baseline is a collection of recommended security settings for your NetBackup domain. For the first time, you configure the security settings as per the recommendation, and use this current posture as your security baseline.

By default, security baseline is not configured.

See [“Security settings to be configured to minimize risk”](#) on page 443.

See [“Set the current posture as security baseline”](#) on page 445.

The security baseline is managed by the NetBackup Administrator or the Security Administrator.

For primary servers that are registered with Cohesity Alta View server, the security baseline is managed by the Cohesity Alta View Administrator.

Recommended values

Recommended values for any given setting represents the possible values of the setting which are considered secure. It is highly recommended to align the current posture of every setting as per the recommended values in order to observe the most secure environment.

Compliance status

If a NetBackup security setting (current posture) does not comply with the security baseline, it is shown in the compliance status as 'Not compliant with the baseline'.

You should review the compliance status and modify the security settings to minimize the risk.

RBAC roles and permissions

To view the **Security configuration risk** card on the NetBackup web UI dashboard, you should have the following roles and permissions:

See [“Configuring RBAC”](#) on page 517.

RBAC roles

Custom

Permissions

View, update global security settings

Security settings to be configured to minimize risk

Configure the following security settings to minimize the security configuration risk.

See [“About security configuration risk”](#) on page 441.

Table 28-1

Security settings	Description
Insecure communication with 8.0 and earlier hosts	This setting determines if insecure communication with 8.0 and earlier hosts is enabled or not. It is recommended that you disable the setting to ensure only the secure communication in the domain.
Security level for certificate deployment	Determines the checks that are performed before the NetBackup CA issues a certificate to a NetBackup host. It is recommended that you set it to High or Very High. See “About NetBackup certificate deployment security levels” on page 475.
Multifactor authentication (MFA)	This setting adds an additional layer of protection in addition to passwords that significantly reduces the risk of malicious access. Enforcing multifactor authentication for all users is recommended. See “Enforce multifactor authentication for all users” on page 469.
Secure data-in-transit encryption (DTE)	This setting determines the global data-in-transit encryption (DTE) mode. It is recommended that you set it to Enforced or Preferred On. See “Configure the global data-in-transit encryption setting” on page 474.
Percent of hosts with DTE enabled	This setting determines the percentage of active hosts in the domain that are participating in DTE.

Table 28-1 (continued)

Security settings	Description
Multi-person authorization (MPA)	This setting ensures that critical actions or decisions are approved by multiple authorized individuals, minimizing the risk of errors, fraud, or misuse of privileges. Enabling this setting is recommended. See “Configure multi-person authorization” on page 457.
Malware detection	This setting determines if malware detection is configured or not. Malware detection scans backup images and detects malware. Configuring malware detection is recommended.
Anomaly detection	This setting detects any unusual deviation in backup job or system attributes and notifies it as an anomaly. Enabling backup and system anomaly detection is recommended. See “Configure backup anomaly detection settings” on page 539. See “Configure system anomaly detection settings” on page 544.
Percent of hosts with service user configured	Measures the percentage of active hosts that are configured to run NetBackup services under a service user account. Having NetBackup services configured to run under a service user (non-privileged user) account is highly recommended. Security configuration risk can be reduced if more hosts are configured to run NetBackup services under service user account. Active primary server, media server, and client hosts are considered for service user configuration.
Percent of encryption-enabled backup storage	This setting identifies the percentage of total active backup storage that is configured to encrypt the data at rest.
Immutable backup storage	This setting identifies if there is at least one active WORM backup storage to be configured. It can either be a storage unit or a tape volume.
Percent of servers with version (primary version) or later	This setting represents the percentage of active hosts (primary and media servers) with NetBackup version later or same as the primary server.
Percent of other hosts with version (primary version) or later	This setting represents the percentage of active hosts (other than primary and media servers) with NetBackup version later or same as the primary server.
CLI access to OS administrator	This setting enables or disables the CLI access for the operating system administrator. It is recommended to disable the setting.
Web UI access to OS administrator	This setting enables or disables web UI access for the operating system administrator. It is recommended to disable the setting.
Client-initiated redirected restores	This setting determines if client-initiated redirected restores are allowed in the domain. It depends on the presence of the <code>No.Restrictions</code> file. It is recommended to remove this file if it exists.

Set the current posture as security baseline

Current posture comprises the current values of NetBackup security settings. It is recommended that you enable all security settings to minimize the security configuration risk.

See [“About security configuration risk”](#) on page 441.

Security baseline is a collection of recommended security settings for your NetBackup domain. For the first time, you configure the security settings as per the recommendation, and use this current posture as your security baseline.

Note: If API is used to set the security baseline, ensure that you provide at least one attribute. Rest of the attributes are set to null. Use the current posture values as security baseline.

At any point in time, a user with the view and update global security settings permissions can set the current configuration posture as the security baseline.

Set the current posture as security baseline

- 1 Enable the security settings to minimize the security configuration risk.
See [“Security settings to be configured to minimize risk”](#) on page 443.
- 2 On the **Global security settings > Overview** tab, click **Set security baseline**.
After the security baseline is set, if a NetBackup security setting is modified and therefore it no longer complies with the security baseline, it is flagged in the compliance status on the **Global security settings > Overview** tab.

Manage security baseline

You can manage or modify the security baseline that you have set earlier.

See [“Set the current posture as security baseline”](#) on page 445.

You can import, export or remove the security baseline. You can also manage the security baseline from the Alta View UI if the associated Alta View server is registered with NetBackup.

Note: The security baseline cannot be modified if it is managed by Alta View administrator. In this case, you can only use the **Export security baseline** option.

To manage security baseline

1 Click **Settings > Global security**.

2 On the **Global security settings > Overview** tab, click **Manage security baseline**.

If you have not set the security baseline earlier, the **Manage security baseline** option is not visible.

3 On the **Manage security baseline** pop-up screen, use one of the following options:

- **Use current posture values as security baseline** - The current posture (security settings values) is set as security baseline.
- **Import security baseline** - An authorized user can set the security baseline by using the import option. This option allows you to provide a JSON file with the security baseline data. The JSON file is created by exporting the security baseline that is set for a NetBackup domain. Select a JSON file with a size that is within the allowed range of 10 KB.
- **Export security baseline** - An authorized user can export the current security baseline into a JSON file. This JSON file can further be used to set the security baseline by providing this file for the import option.
- **Remove the current security baseline** - An authorized user can delete the security baseline. If a security baseline is set by an Alta View user, it can be modified only by the Alta View user.
If security baseline is deleted, notifications are not generated for any changes in the security risk.

Manage security baseline from Alta View UI

If a security baseline is set by an Alta View user, it cannot be modified or deleted by NetBackup users. Only Alta View users can modify the security baseline.

If the baseline is set by an Alta View user, the **Manage security baseline** option is disabled for NetBackup users.

See [“Manage security baseline”](#) on page 445.

Configuring multi-person authorization

This chapter includes the following topics:

- [About multi-person authorization](#)
- [Workflow to configure multi-person authorization for NetBackup operations](#)
- [RBAC roles and permissions for multi-person authorization](#)
- [Multi-person authorization process with respect to roles](#)
- [NetBackup operations that need multi-person authorization](#)
- [Configure multi-person authorization](#)
- [View multi-person authorization tickets](#)
- [Manage multi-person authorization tickets](#)
- [Add exempted users](#)
- [Schedule expiration and purging of multi-person authorization tickets](#)
- [Disable multi-person authorization](#)

About multi-person authorization

NetBackup Security Administrator can configure multi-person authorization that helps protect primary servers from an undesirable or a malicious act, in a proactive manner. Multi-person authorization ensures that a second authorized user approves actions before they are performed.

To configure multi-person authorization in NetBackup, you need to have two users: one is the requester and the other is the approver.

A requester cannot be an approver of their own tickets.

Support information

- Multi-person authorization is not supported in a domain where NetBackup Access Control (NBAC) is enabled.
- Multi-person authorization is not supported for catalog maintenance operations by certain database agents.

As part of the database catalog synchronization, the database may initiate an image expiration request through command-line or other interfaces to the NetBackup catalog, which does not generate multi-person authorization ticket. To prevent the direct expiration of backup images by database agents see the 'About preventing the direct expiration of backup images' topic in the [NetBackup for Oracle Administrator's Guide](#).

Terminology

- Ticket - Ticket is a multi-person authorization request to perform a critical operation.
- Requester - A requester is a user who wants to perform a critical operation that requires multi-person authorization.
- Approver - An approver is an individual who reviews and allows an operation that requires multi-person authorization by approving a ticket.
- Exempted user - An exempted user is not required to go through the multi-person authorization workflow. This user must only be used to perform critical operations like image expiration and image hold removal.

For additional security, it is recommended that there are no exempted users.

Command-line options that need multi-person authorization

The following operations and the associated command-line options need multi-person authorization:

- Expiring images expiration:
 - `bpexpdate`
 - `nbdecommission`
 - `bpimage -deleteCopy`
- Removing image hold:

- `nbholdutil -delete`
- Modifying global security settings:
 - `nbcertcmd -setsecconfig`
 - `nbseccmd -setsecurityconfig`
- Managing encryption key
 - `nbkmscmd`
 - `nbkmsutil`

For more information on commands, see the [NetBackup Command Reference Guide](#).

Multi-person authorization is supported for the following commands that are run with the `nbcmdrun` command:

- `bpplcatdrinfo`
- `bpplclients`
- `bppldelete`
- `bpplsched`
- `bpplininclude`
- `bpplininfo -set`
- `bpplsched`
- `bpplschedrep`
- `bpplschedwin`
- `bpolicynew`

Multi-person authorization in a NetBackup and Alta View setup

If an Alta View user has requested a NetBackup operation that needs multi-person authorization, on a registered primary server, multi-person authorization must be enabled in Alta View. Else, NetBackup rejects this Alta View request and the respective user operation fails.

Workflow to configure multi-person authorization for NetBackup operations

Here are the high-level steps to configure multi-person authorization for NetBackup operations:

Table 29-1

Step	Description
Step 1	Identify critical NetBackup operations that require multi-person authorization. See “NetBackup operations that need multi-person authorization” on page 455.
Step 2	Identify the approvers who can approve requests or multi-person authorization tickets.
Step 3	Assign the Default multi-person authorization approver RBAC role to the approvers. See “RBAC roles and permissions for multi-person authorization” on page 451.
Step 4	Configure multi-person authorization using the NetBackup web UI. See “Configure multi-person authorization” on page 457.
Step 5	When a user or a requester tries to perform an operation that requires multi-person authorization (for example, expiring an image), a ticket is generated. Initially, the ticket is in the pending state.
Step 6	The ticket is visible to all multi-person authorization approvers in the NetBackup web UI where they can review the ticket information and approve or reject the ticket.
Step 7	When the approver approves or rejects the ticket, the requester is notified. If the ticket is approved, the associated operation is executed. Note: For API key operations, the requester needs to execute the operation using the web UI after the ticket is approved.

Multi-person authorization configuration begins when the Administrator or the Security Administrator enables critical operations that require multi-person authorization and specifies other settings like expiration period and purge period.

A multi-person authorization configuration ticket is generated. After the approver approves the ticket, multi-person authorization configuration comes into effect.

Initial multi-person authorization configuration

Configuring multi-person authorization for the first time involves adding users to the Default Multi-Person Authorization Approver role. To start using the multi-person authorization for additional data security, the Security Administrator must enable the multi-person authorization for critical pre-defined operations that require an additional approval from a user with the Default Multi-Person Authorization Approver role.

Initially, the Security Administrator should configure multi-person authorization that results into a multi-person authorization ticket. After the approver approves the ticket, multi-person authorization becomes mandatory for the specified NetBackup operation (such as image expiry). The Administrator or Security Administrator can add users to the Default Multi-Person Authorization Approver role at any point in time.

RBAC roles and permissions for multi-person authorization

Multi-person authorization configuration requires the users to be assigned to the following RBAC roles:

- Administrator
- Default Security Administrator
- Default Multi-Person Authorization Approver

Users with these RBAC roles should have the following permissions.

Table 29-2

RBAC role	Permissions
Administrator	View, update multi-person authorization configuration, and delegate the configuration permissions to other users. View, update tickets, and delegate ticket permissions to other users.
Default Security Administrator	View, update multi-person authorization configuration, and delegate the configuration permissions to other users.

Table 29-2 *(continued)*

RBAC role	Permissions
Default Multi-person Authorization Approver	View and update tickets.
Default Operator	View all NetBackup entities.

Multi-person authorization process with respect to roles

Users can be requesters and approvers at the same time, however they cannot approve their own tickets.

The multi-person authorization process flow with respect to roles is as follows:

Table 29-3

Component	Description
Multi-person authorization ticket	When a requester performs a critical NetBackup operation that is protected by multi-person authorization, a ticket is generated that requires an approval from the approver before a specific action can be executed. This ticket is used within NetBackup to ensure that critical actions undergo thorough review process by multiple people before they are executed. The following sample flow is for the image expiry operation that requires multi-person authorization: 1 A requester expires an image using the NetBackup web UI. 2 A ticket is created. 3 The ticket is pending for approval. 4 Approvers review the ticket. 5 Approvers either approve or reject the ticket. 6 After the approval, the ticket is scheduled by NetBackup and finally marked Done after the execution. 7 The ticket activity log, request, and response details can be viewed by the approver or the requester using the web UI, on the Ticket details page. 8 A ticket is expired after it ages beyond the expiration period. Such tickets cannot be approved unless they are renewed by the Requester. 9 Tickets in the Done, Rejected, Expired, and Canceled states are purged when no action is performed on them for the specified purge period in days.

Table 29-3 (continued)

Component	Description
Requester role	1 A requester is a user who initiates an operation that requires multi-person authorization. 2 A ticket is created for the operation if the user is not in the exempted users' list. 3 The ticket requires an approval from an approver before the operation is performed. 4 A requester is not allowed to self approve even if the requester is also an approver, an Administrator, or a Security Administrator. 5 Once the ticket is created it is in the Pending state. 6 The requester can cancel a ticket only if it is in the Pending state. 7 If the ticket ages beyond the expiry period, the ticket is moved to the Expired state. 8 Only the requester can renew such tickets. A new expiry period is calculated for the renewed ticket based on the configuration settings multi-person authorization.
Approver role	1 An approver is an authorized individual who reviews and provides approval for tickets. 2 The approver evaluates the details of the ticket and either approves or rejects the ticket based on the assessment. 3 After the approval, the ticket is scheduled for execution. 4 To be an approver, the user should have RBAC permissions like Update Ticket, View Ticket or the user should have the Default Multi-Person Authorization Approver role. 5 When a ticket is in the Pending State, it can be approved or rejected.

Table 29-3 (continued)

Component	Description
Exempted users	1 An exempted user is an individual who does not need multi-person authorization for operations except the following: ■ To modify multi-person authorization configuration ■ To modify security properties 2 User groups cannot be exempted. 3 This eliminates the necessity for any approvals, however it must be used with caution. 4 If the exempted user account is hacked, the multi-person authorization process can be of no use as it is bypassed for this user. 5 For example, if user1 is an exempted user and she attempts to expire an image (an operation that needs multi-person authorization), the image expires without ticket generation and additional approvals.

NetBackup operations that need multi-person authorization

The following operations require multi-person authorization and therefore a ticket is generated for these operations:

- Configuring multi-person authorization
- Enabling and disabling operations that require multi-person authorization
- Adding exempted users
- Changing any multi-person authorization settings
- Expiring images
- Updating image expiration time
- Changing the MSDP WORM configuration
- Removing the MSDP WORM retention lock
- Removing hold applied on the images
- Updating CLI expiration period
- Adding, updating, and deleting an API key
- Adding, updating, and deleting KMS configuration, keys, and key groups

- Adding, updating, deleting malware scan host
- Adding, updating, deleting, copying backup and deployment policies
- Updating the following global security settings:
 - Enabling and disabling NetBackup host communication with insecure hosts
 - Adding host aliases with or without NetBackup administrator's approval
 - Setting automatic deployment of certificates on a host
 - Enabling and disabling CAC/PIV authentication
 - Setting values for CAC/PIV certificate mapping attribute
 - Setting the value of the CAC/PIV certificate mapping attribute that is used to perform a search in active directory
 - Setting the value of the CAC/PIV certificate mapping attribute that is used to perform a search in LDAP directory
 - Enabling and disabling AD/LDAP domain mapping
 - Setting the value of the domain name that is used for user look-ups in active directory or LDAP
 - Setting the value of the OCSP URI that is used for certificate revocation checks with respect to CAC/PIV authentication
 - Enabling and disabling the data-in-transit encryption (DTE)
 - Setting unique identifier for external certificates
 - Allowing or disallowing the NetBackup web UI access to Operating System Administrators
 - Allowing or disallowing the default CLI access to OS administrators
 - Pausing client protection
 - Pausing client image expiration
 - Enabling and disabling TLS session resumption
 - Enabling and disabling rule engine for anomaly detection
 - Changing multifactor authentication configuration settings
 - Setting audit retention period for audit report

Even if multi-person authorization is configured for image expiry, the following operations do not require multi-person authorization:

- Changing values for image retention level
- Modifying retention levels in policy and SLP

- Canceling incomplete SLPs using the `nbstlutil` command:
Refer to the *NetBackup Commands Reference Guide*.

Configure multi-person authorization

The configuration of multi-person authorization for NetBackup operations is supported only from the NetBackup web UI. A user with the Administrator or the Security Administrator role can configure multi-person authorization for critical NetBackup operations.

To configure multi-person authorization for NetBackup operations

- 1 On the left, select **Security > Multi-person authorization**.
- 2 At the top right, select the option **Configure multi-person authorization**.
- 3 Go to **Operations for multi-person authorization**. Then select **Edit**.
- 4 Select all or any of the following critical operations for which you want to configure multi-person authorization.
 - Images
 - Image expiry
 - Remove image hold
 - Security
 - Global security settings
 - Encryption key management
 - API keys

Note: If multi-person authorization is enabled for API key operations, a ticket is generated. After the multi-person authorization ticket is approved, the user needs to execute the ticket using the **Execute ticket** option in the NetBackup web UI and then the required API key operation is executed.

For NetBackup releases earlier than 10.5, if multi-person authorization is enabled, you cannot perform API key operations.

- Malware scan host management

Note: Starting with NetBackup 11.0, if multi-person authorization is enabled for this operation, a ticket is generated.

- Protection
 - Policy management

Note: Starting with NetBackup 11.0, if multi-person authorization is enabled for this operation, a ticket is generated. If there is an existing ticket for a policy that is pending for approval, you cannot perform any operations on the same policy until the existing ticket is either approved, rejected, canceled, or expired.

- MSDP WORM
 - WORM retention lock removal
 - WORM configuration change

- 5 Select **Save**.
- 6 Configure the users to be exempted from multi-person authorization.
Starting with 11.0, you can exempt users from multi-person authorization for specific operation and action for specific time frame.
- 7 Configure email notifications that you want to send to approvers and other email recipients. You can specify additional email recipients to whom you want to send email notifications.

Click **Configure** to configure the SMTP server for emails.
- 8 Go to **Schedules**. Then select **Edit**.
- 9 Specify when you want to expire and purge the multi-person authorization tickets.
- 10 Select **Save**.
- 11 Select **Configure**.

View multi-person authorization tickets

Users can view their own multi-person authorization tickets.

- ◆ On the left pane, click **Security > Multi-person authorization**. The list of multi-person authorization tickets is displayed.

Select the ticket ID to see more details.

Manage multi-person authorization tickets

Users with the approver role can approve or reject the multi-person authorization tickets.

To manage multi-person authorization tickets

- 1 Sign into the NetBackup web UI.
- 2 On the left pane, click **Security > Multi-person authorization**. The list of multi-person authorization tickets is displayed.
- 3 Select the ticket ID to view the request details.
- 4 Select **Review ticket** and enter the respective ticket ID to approve or reject the ticket.
- 5 Add comments and click **Approve** or **Reject**.

Add exempted users

Your organization may need to exempt certain users from multi-person authorization so that operations like image expiry and image hold removal can proceed without second-level approval.

Add such users to the exempted users' list.

Note: User groups cannot be added to the exempted list. Only individual users can be exempted.

The exempted users also need to go through the multi-person authorization workflow for the following operations:

- Modifying multi-person authorization configuration
- Modifying global security settings
- Modifying risk engine-based anomaly detection configuration

An exempted user is generally an automation user or a script that does not require multi-person authorization. By default, multi-person authorization configuration does not have exempted users and that is a recommended security setting.

To add exempted users

- 1 Sign into the NetBackup web UI.
- 2 On the left pane, select **Security > Multi-person authorization**.
- 3 On the top right, select **Configure multi-person authorization**.

- 4 In the **Exempted users** section, select the **Add** button.
- 5 Specify the name of the user whom you want to exempt from the multi-person authorization process.
- 6 Select **Add to list** and then **Save**.
- 7 Select **Save**.

Schedule expiration and purging of multi-person authorization tickets

Expiration period is configurable option defines the duration for which a multi-person authorization ticket can be in the Pending state. A ticket expires if it is in the Pending state for more than the configured expiry period.

For multi-person authorization configuration, expiration period can vary from minimum 24 hours to 168 hours. By default, tickets expire after 72 hours.

Purge period is a configurable option defines the duration for which a ticket resides in the tickets database. Purging a ticket ensures that the database does not grow exponentially. Purge period can vary from minimum 3 days to 30 days.

By default, tickets purge after 72 hours. All the Done, Expired, Rejected, and Canceled tickets are purged after the given purge period.

To schedule expiration and purging of tickets

- 1 On the left pane, click **Security > Multi-person authorization**.
- 2 On the top right, click **Configure multi-person authorization**.
- 3 In the **Schedules** section, click **Edit**.
- 4 Specify the expiration period (in hours) for the **Expire ticket after** option.
Specify the purge period (in days) for the **Purge ticket after** option.
- 5 Select **Save**.
- 6 Select **Save**.

Disable multi-person authorization

In certain cases, you may need to temporarily disable multi-person authorization for the associated operations.

To disable multi-person authorization for all the associated operations, run the following command after `bpnbat -login -loginType WEB` using the root or Administrator account.

```
nbseccmd -disableMPA
```

You can disable multi-person authorization for a specific operation using the NetBackup web UI

To disable multi-person authorization for a specific operation

- 1** On the left pane, click **Security > Multi-person authorization**.
- 2** On the top right, click **Configure multi-person authorization**.
- 3** In the **Operations for multi-person authorization** section, click **Edit**.
- 4** Clear the check box for the operation for which you want to disable multi-person authorization.
- 5** Select **Save**.
- 6** Select **Save**.

This generates a ticket that is shown on the ticket details page with the operation name as MPA Configuration.

Multi-person authorization will be disabled for the associated operation only after the approval of the respective ticket.

Managing user sessions

This chapter includes the following topics:

- [Terminate a NetBackup user session](#)
- [Unlock a NetBackup user](#)
- [Configure when idle sessions should time out](#)
- [Configure the maximum of concurrent user sessions](#)
- [Configure the maximum of failed sign-in attempts](#)
- [Display a banner to users when they sign in](#)

Terminate a NetBackup user session

For security or maintenance purposes, you can terminate one or more NetBackup user sessions. To configure NetBackup to automatically terminate any idle user sessions, see the following topic.

See [“Configure when idle sessions should time out”](#) on page 464.

Note: Changes to a user’s roles are not immediately reflected in the web UI. An administrator must terminate the active user session before any changes take effect. Or, the user must sign out and sign in again.

To sign out a user session

- 1 On the left, click **Security > User sessions**.
- 2 At the top right, click **User account settings**.
- 3 Go to the **Active sessions** tab.

- 4 Select the user session that you want to sign out.
- 5 Select **Terminate session**.

To sign out all user sessions

- 1 On the left, click **Security > User sessions**.
- 2 At the top right, click **User account settings**.
- 3 Go to the **Active sessions** tab.
- 4 Select **Terminate all sessions**.

Unlock a NetBackup user

You can view the user accounts that are currently locked out of NetBackup and unlock one or more users.

By default a user's account only remains locked for 24 hours. You can change this time by adjusting the **User sessions > User account settings > User account lockout** setting.

See [“Configure the maximum of failed sign-in attempts”](#) on page 464.

To unlock out a locked user account

- 1 On the left, click **Security > User sessions**.
- 2 At the top right, click **User account settings**.
- 3 Go to the **Locked users** tab.
- 4 Select the user account that you want to unlock.
- 5 Select **Unlock**.

To unlock all locked user accounts

- 1 On the left, click **Security > User sessions**.
- 2 At the top right, click **User account settings**.
- 3 Go to the **Locked users** tab.
- 4 Select **Unlock all users**.

Configure when idle sessions should time out

You can customize when user sessions should time out and a user is automatically signed out. The setting you choose is applied to the NetBackup web UI. To configure this setting from the command line, use `nbsetconfig` to set the `GUI_IDLE_TIMEOUT` option.

To configure when idle sessions should time out

- 1 On the left, click **Security > User sessions**.
- 2 At the top right, click **User account settings**.
- 3 Turn on **Session idle timeout** and click **Edit**.
- 4 Select the number of minutes and click **Save**.

For active users, the updates are applied the next time the user signs in.

Configure the maximum of concurrent user sessions

This setting limits the number of concurrent API sessions that a user can have active. This setting does not apply to API key sessions or to other applications like the NetBackup Backup, Archive, and Restore interface.

To configure this setting from the command line, use `nbsetconfig` to set the `GUI_MAX_CONCURRENT_SESSIONS` option.

To configure the maximum of concurrent user sessions

- 1 On the left, click **Security > User sessions**.
- 2 At the top right, click **User account settings**.
- 3 Turn on **Maximum concurrent sessions** and click **Edit**.
- 4 Select the **Number of concurrent sessions per user** and click **Save**.

For active users, the updates are applied the next time the user signs in.

Configure the maximum of failed sign-in attempts

You can automatically lock a user account if the user exceeds a maximum number of failed sign-in attempts. The user account remains locked until the account lockout period passes.

If there is an immediate need to access NetBackup, the administrator can unlock the account.

See [“Unlock a NetBackup user”](#) on page 463.

You can customize the maximum number of NetBackup failed sign-in attempts. The setting you choose applies only to the NetBackup web UI. To configure this setting from the command line, use `nbsetconfig` to set the `GUI_MAX_LOGIN_ATTEMPTS` and `GUI_ACCOUNT_LOCKOUT_DURATION` options.

To configure the maximum of failed sign-in attempts

- 1 On the left, click **Security > User sessions**.
- 2 At the top right, click **User account settings**.
- 3 Turn on **User account lockout** and click **Edit**.
- 4 Select the number of failed sign-in attempts that you want to allow before an account is locked.
- 5 To unlock a locked account after a period of time, select the number of minutes for **Unlock locked accounts after**.
- 6 Select **Save**.

For active users, the updates are applied the next time the user signs in.

Display a banner to users when they sign in

You can configure a sign-in banner that displays each time that any user signs in to the NetBackup web UI. A different banner can be configured for any primary server. This banner can also require the user to agree to the terms of service before the user signs in.

To display a banner to users when they sign in

- 1 On the left, click **Security > User sessions**.
- 2 At the top right, click **User account settings**.
- 3 Turn on **Sign-in banner configuration** and click **Edit**.
- 4 Enter the text you want to use for the heading and the body of the message.
- 5 If you want to require the user to agree to the terms of service, select **Include "Agree" and "Disagree" buttons on the sign-in banner**.
- 6 Select **Save**.

For active users, the updates are applied the next time the user signs in.

To remove the sign-in banner

- 1** On the left, click **Security > User sessions**.
- 2** At the top right, click **User account settings**.
- 3** Turn off **Sign-in banner configuration**
- 4** Select **Save**.

For active users, the updates are applied the next time the user signs in.

Configuring multifactor authentication

This chapter includes the following topics:

- [About multifactor authentication](#)
- [Configure multifactor authentication for your user account](#)
- [Disable multifactor authentication for your user account](#)
- [Enforce multifactor authentication for all users](#)
- [Configure multifactor authentication for your user account when it is enforced in the domain](#)
- [Reset multifactor authentication for a user](#)

About multifactor authentication

Multifactor authentication is a multiple-step account login process that requires you to enter a 6-digit one-time password along with your password.

It is strongly recommended that you configure multifactor authentication to protect the security of your environment.

Note: User logins that are based on the following authentication types do not support multifactor authentication: SAML, smart card, and API keys.

See [“Configure multifactor authentication for your user account”](#) on page 468.

If multifactor authentication is configured, you may need to reauthenticate yourself by entering the one-time password that you see in the authenticator application on your smart device before you perform the following operations:

- Manage the global security settings for the primary server
- Adding an API key
See [“Add an API key or view your API key details”](#) on page 489.

If multifactor authentication is enforced in the NetBackup domain, all users must configure multifactor authentication for their user accounts for successful sign-in.

See [“Configure multifactor authentication for your user account when it is enforced in the domain”](#) on page 470.

Configure multifactor authentication for your user account

For enhanced security, you can configure multifactor authentication for your user account. You must first install and configure authenticator application on your smart device that provides you with the one-time password.

Configuring multifactor authentication in NetBackup does not require internet connectivity on your smart device.

If the NetBackup administrator has enforced multifactor authentication in the NetBackup domain, you must configure it for your user account for successful sign-in.

See [“Disable multifactor authentication for your user account”](#) on page 469.

To configure multifactor authentication for your user account

- 1 On the top right, click the profile icon and click **Configure multifactor authentication**.
- 2 On the **Configure multifactor authentication** screen, click **Configure**.
- 3 On the next screen, follow the given steps.

Install and configure authenticator application on your smart device. It generates one-time password and sends it on your smart device.

[Supported authenticator applications](#)

- 4 Scan the QR code with the authenticator application or enter the key manually.

- 5 Enter the one-time password that you see in the authenticator application on your smart device.
- 6 Select **Configure**.

At the time of next sign-in, you need to enter the one-time password along with the username and password.

Disable multifactor authentication for your user account

If multifactor authentication is not enforced, you can disable it for your user account. However, it is strongly recommended that you configure multifactor authentication to protect the security of your account.

See [“Configure multifactor authentication for your user account”](#) on page 468.

To disable multifactor authentication for your user account

- 1 On the top right, click the profile icon and select **Configure multifactor authentication**.
- 2 If you have already configured multifactor authentication for your user account, you can see the **Disable** option.
- 3 Select **Disable**.
- 4 Enter the one-time password and click **Confirm**.

Enforce multifactor authentication for all users

Only the NetBackup administrator can enforce multifactor authentication for all NetBackup users.

To enforce multifactor authentication for all users

- 1 On the top right, click **Settings > Global security**.
- 2 On the **Security controls** tab, turn on **Enforce multifactor authentication**.
Select **Confirm** to enforce multifactor authentication for all NetBackup users.
Notify all users that they must configure multifactor authentication for their user accounts to be able to successfully sign in.

See [“Configure multifactor authentication for your user account”](#) on page 468.

Configure multifactor authentication for your user account when it is enforced in the domain

After multifactor authentication is enforced in the domain, you must configure it for your user account if you have not already configured it. If you do not configure multifactor authentication for your account after the enforcement, you cannot sign-in.

To configure multifactor authentication after the enforcement

- 1 Open a web browser and go to the following URL.
`https://primaryserver/webui/login`
 The *primaryserver* is the host name or IP address of the NetBackup primary server that you want to sign in to.
- 2 Go to the NetBackup sign-in screen.
- 3 Enter the **Username** and **Password**.
 See “Sign in to the NetBackup web UI” on page 32.
- 4 Select **Sign in**. The **Configure multifactor authentication** screen is displayed.
- 5 On the next screen, follow the given steps.
 Install and configure an authenticator application on your smart device. It generates a one-time password and sends it to your smart device.
[Supported authenticator applications](#)
- 6 Scan the QR code with the authenticator application or enter the key manually.
- 7 Enter the one-time password that you see in the authenticator application on your smart device.
- 8 Select **Configure**.
 Successful configuration takes you back to the sign-in screen.
 Enter the username, password, and one-time password for successful sign-in.

Reset multifactor authentication for a user

Only the NetBackup administrator can reset multifactor authentication for other NetBackup users.

To reset multifactor authentication for a NetBackup user

- 1 On the top right, click **Settings > Global security**.
- 2 Go to the **Security controls** tab.

- 3** Locate the section **Reset multifactor authentication for a user**. Then select **Reset**.
- 4** Select the user for whom you want to reset multifactor authentication.
- 5** Select **Reset**.
- 6** At the prompt, enter the one-time password and select **Confirm**.

Managing the global security settings for the primary server

This chapter includes the following topics:

- [View the Certificate authority for secure communication](#)
- [Disable communication with NetBackup 8.0 and earlier hosts](#)
- [Disable automatic mapping of NetBackup host names](#)
- [Configure the global data-in-transit encryption setting](#)
- [About NetBackup certificate deployment security levels](#)
- [Select a security level for NetBackup certificate deployment](#)
- [About TLS session resumption](#)
- [Set a passphrase for disaster recovery](#)
- [Validate the disaster recovery package passphrase](#)
- [About trusted primary servers](#)
- [Configure the audit retention period](#)

View the Certificate authority for secure communication

In the global security settings, the **Certificate authority** information indicates the type certificate authorities that the NetBackup domain supports.

NetBackup hosts in the domain can use certificates as follows:

- **NetBackup certificates.**
By default, NetBackup certificates are deployed on the primary server and its clients.
- **External certificates.**
You can configure NetBackup to only communicate with the hosts that use an external certificate. This configuration requires that a host is upgraded to 8.2 or later and has an external certificate that is installed and enrolled. In this case, NetBackup does not communicate with any hosts that use NetBackup certificates. However, you can enable **Allow communication with NetBackup 8.0 and earlier hosts** to communicate with any hosts that use NetBackup 8.0 or earlier.
- **Both NetBackup certificates and external certificates.**
With this configuration, NetBackup communicates with the hosts that use a NetBackup certificate or an external certificate. If a host has both types of certificates, NetBackup uses the external certificate for communication.

To view the certificate authorities that a NetBackup domain supports

- 1 Open the NetBackup web UI.
- 2 At the top right, select **Settings > Global security**.
- 3 Go to the **Secure communication** tab.
- 4 Go to the **Certificate Authority** section. This section lists the CAs that NetBackup supports.

Disable communication with NetBackup 8.0 and earlier hosts

NetBackup allows communication with NetBackup 8.0 and earlier hosts that are present in the environment. However, this communication is insecure. For increased security, upgrade all your hosts to the current NetBackup version and disable this setting. This action ensures that only secure communication is possible between NetBackup hosts. If you use Auto Image Replication (A.I.R.), you must upgrade the trusted primary server for image replication to NetBackup 8.1 or later.

To disable communication with NetBackup 8.0 and earlier hosts

- 1 At the top right, select **Security > Global security**.
- 2 Select the **Secure communications** tab.
- 3 Turn off **Enable communication with NetBackup 8.0 and earlier hosts**.
- 4 Select **Save**.

Disable automatic mapping of NetBackup host names

For successful communication between NetBackup hosts, all relevant host names and IP addresses need to be mapped to the respective host IDs. Use the **Automatically map host names to their NetBackup host ID** option to automatically map the host ID to the respective host names (and IP addresses). Or, disable it to allow the NetBackup security administrator to manually verify the mappings before approving them.

To disable automatic mapping of NetBackup host names

- 1 At the top right, select **Settings > Global security**.
- 2 Select the **Secure communications** tab.
- 3 Turn off **Automatically map NetBackup host ID to host names**.
- 4 Select **Save**.

Configure the global data-in-transit encryption setting

To configure the data-in-transit encryption (DTE) in your NetBackup environment, you need to first set the global DTE configuration setting (or global DTE mode) and then the client DTE mode.

Data-in-transit encryption decision for various NetBackup operations is carried out based on the global DTE mode, the client DTE mode, and the image DTE mode.

The supported values for the global DTE mode are as follows:

- **Preferred Off**: Specifies that the data-in-transit encryption is disabled in the NetBackup domain. This setting can be overridden by the NetBackup client setting.
- **Preferred On**: Specifies that the data-in-transit encryption is enabled only for NetBackup 9.1 and later clients.

In case of fresh NetBackup installation, the global DTE mode is set to `Preferred On` by default.

In case of NetBackup upgrade, the previous setting is retained.

This setting can be overridden by the NetBackup client setting.

- **Enforced:** Specifies that the data-in-transit encryption is enforced if the NetBackup client setting is either 'Automatic' or 'On'. With this option selected, jobs fail for the NetBackup clients that have the data-in-transit encryption set to 'Off' and for the hosts earlier than 9.1.

Note: By default, the DTE mode for 9.1 clients is set to `off` and for 10.0 and later clients, it is set to `Automatic`.

RESTful API to be used for the global DTE configuration:

- GET - `/security/properties`
- POST - `/security/properties`

To set or view the global DTE mode using the NetBackup web UI

- 1 At the top right, select **Security > Global security**.
- 2 On the **Secure communication** tab, select one of the following global DTE settings:
 - Preferred Off
 - Preferred On
 - Enforced

About NetBackup certificate deployment security levels

Security levels for certificate deployment are specific to NetBackup CA-signed certificates. If the NetBackup web server is not configured to use NetBackup certificates for secure communication, the security levels cannot be accessed.

The NetBackup certificate deployment level determines the checks that are performed before the NetBackup CA issues a certificate to a NetBackup host. It also determines how frequently the NetBackup Certificate Revocation List (CRL) is refreshed on the host.

NetBackup certificates are deployed on hosts during installation (after the host administrator confirms the primary server fingerprint) or with the `nbcertcmd`

command. Choose a deployment level that corresponds to the security requirements of your NetBackup environment.

Note: During NetBackup certificate deployment on a NAT client, you must provide an authorization token irrespective of the certificate deployment security level that is set on the primary server. This is because, the primary server cannot resolve the host name to the IP address from which the request is sent.

For more information about NAT support in NetBackup, refer to the [NetBackup Administrator's Guide, Volume I](#).

Table 32-1 Description of NetBackup certificate deployment security levels

Security level	Description	CRL refresh
Very High	An authorization token is required for every new NetBackup certificate request.	The CRL that is present on the host is refreshed every hour.

Table 32-1 Description of NetBackup certificate deployment security levels
(continued)

Security level	Description	CRL refresh
High (default)	No authorization token is required if the host is known to the primary server. A host is considered to be known to the primary server if the host can be found in the following entities: 1 The host is listed for any of the following options in the NetBackup configuration file (Windows registry or the <code>bp.conf</code> file on UNIX): ■ APP_PROXY_SERVER ■ DISK_CLIENT ■ ENTERPRISE_VAULT_REDIRECT_ALLOWED ■ MEDIA_SERVER ■ NDMP_CLIENT ■ SERVER ■ SPS_REDIRECT_ALLOWED ■ TRUSTED_MASTER ■ VM_PROXY_SERVER ■ MSDP_SERVER For more details on the NetBackup configuration options, refer to the NetBackup Administrator's Guide, Volume I. 2 The host is listed as a client name in the <code>altnames</code> file (ALTNAMEESDB_PATH). 3 The host appears in the EMM database of the primary server. 4 At least one catalog image of the client exists. The image must not be older than 6 months. 5 The client is listed in at least one backup policy. 6 The client is a legacy client. This is a client that was added using the Client Attributes host properties.	The CRL that is present on the host is refreshed every 4 hours.
Medium	The certificates are issued without an authorization token if the primary server can resolve the host name to the IP address from which the request was originated.	The CRL that is present on the host is refreshed every 8 hours.

Select a security level for NetBackup certificate deployment

NetBackup offers several security levels for the NetBackup certificate deployment. The security level determines what security checks the NetBackup certificate authority (CA) performs before it issues a certificate to a NetBackup host. The level also determines how frequently the Certificate Revocation List (CRL) for the NetBackup CA is refreshed on the host.

More details are available for security levels, NetBackup certificate deployment, and the NetBackup CRL:

- See [“About NetBackup certificate deployment security levels”](#) on page 475.
- See the [NetBackup Security and Encryption Guide](#).

To select a security level for NetBackup certificate deployment

- 1 At the top, select **Settings > Global security**.
- 2 Select the **Secure communication** tab.
- 3 For **Security level for certificate deployment**, select a security level.

If you choose to use NetBackup certificates, they are deployed on hosts during installation, after the host’s administrator confirms the primary server fingerprint. The security level determines if an authorization token is required or not for a host.

Very high	NetBackup requires an authorization token for every new NetBackup certificate request.
High (Default)	NetBackup does not require an authorization token if the host is known to the primary server. Known means that the host appears in a NetBackup configuration file, the EMM database, a backup policy, or the host is a legacy client.
Medium	NetBackup issues NetBackup certificates without an authorization token if the primary server can resolve the host name to the IP address from which the request was originated.

- 4 Select **Save**.

About TLS session resumption

NetBackup uses TLS (Transport Layer Security) to secure communications between NetBackup hosts and is enabled by default. Each new TCP connection between

NetBackup hosts must perform a TLS handshake and verify the peer identity before NetBackup sends traffic across that connection.

TLS session resumption is an open standards optimization that allows a TLS client and server to reuse a secure session that is generated during a previous connection. Reusing a secure session allows NetBackup to use a streamlined handshake instead of a full handshake. Performing this action reduces both the host CPU and time that is required to establish the new connection.

TLS version 1.2 reduces forward security for the interval between full handshakes. To limit this window while still benefitting from session reuse, NetBackup allows global configuration of the maximum interval between full TLS handshakes.

To use the options for **TLS session resumption**, navigate to **Settings > Global security > Secure communication**. You can use the **Perform full handshake every** option to set the security level as follows:

- **Default for current security level** – If you use this option, NetBackup defaults to the security setting as follows:
 - Very high - 10 minutes
 - High - 30 minutes
 - Medium - 60 minutes
- **Custom (overrides the security level settings)** - The value of this interval can be configured at a minute granularity, within the range of 1 minute to 720 minutes.

The TLS 1.3 session ticket lifetime is same as the interval that is mentioned earlier. However, the TLS 1.3 session ticket is used only once.

Note: If strict forward security is a concern, NetBackup also allows session resumption to be globally disabled.

Note: This feature currently only applies to NBBCA. ECA to be supported in a future release.

Set a passphrase for disaster recovery

During a catalog backup, NetBackup creates a disaster recovery package and encrypts the backup with a passphrase that you set. The constraints for the passphrase can be changed with the NetBackup APIs or the CLIs (`nbseccmd -setpassphraseconstraints`).

See the information for disaster recovery settings in the [NetBackup Security and Encryption Guide](#).

To set a passphrase for disaster recovery

- 1 At the top, click **Settings > Global security**.
- 2 Go to the **Disaster recovery** tab.
- 3 Enter and confirm a passphrase.

Note: The passphrase should meet any additional constraints that you may have set. You can verify the additional constraints using the `nbseccmd` command or the passphrase-constraints web API.

- 4 Select **Save**.

Validate the disaster recovery package passphrase

Cohesity strongly recommends that you validate the disaster recovery (DR) package passphrase every 30 days. It helps you recall the passphrase when you need to recover the primary server identity using the DR package.

If validation fails, the specified passphrase is set as your DR package passphrase.

To validate the disaster recovery package passphrase

- 1 At the top right, select **Settings > Global security**.
- 2 Go to the **Disaster recovery** tab.
- 3 Locate the **Passphrase Validation** section and select **Validate**.
- 4 Enter and confirm the passphrase that you set earlier.
- 5 Select the check box **Notify me to validate passphrase every 30 days**. This option is recommended.
- 6 Select **Validate**.

About trusted primary servers

A trust relationship between NetBackup domains lets you do the following:

- Select specific domains as a target for replication. This type of Auto Image Replication is known as targeted A.I.R.

Without a trust relationship, NetBackup replicates to all defined target storage servers. A trust relationship is optional for Media Server Deduplication Pool and PureDisk Deduplication Pool as a target storage. To use a Cloud Catalyst storage server, a trust relationship is required.

- Include usage reporting for multiple primary servers.

Primary servers can use a NetBackup certificate authority (CA) certificate or an external CA certificate. NetBackup determines the CAs used by the source and the target domains and selects the appropriate CA to use for communication between the servers. If the target primary server is configured for both CA types, NetBackup prompts you to select the CA that you want to use. To establish trust with a remote primary server using the NetBackup CA, the current primary and the remote primary must have NetBackup version 8.1 or later. To establish trust with a remote primary server using an external CA, the current primary and the remote primary must have NetBackup version 8.2 or later.

About the certificate to use to add a trusted primary server

A source or a target primary server may use NetBackup CA-signed certificates (host ID-based certificates) or external CA-signed certificates.

For more information on NetBackup host ID-based certificates and external CA support, refer to the [NetBackup Security and Encryption Guide](#).

To establish trust between source and target primary servers, NetBackup verifies the following:

Can the source primary server establish trust using an external CA-signed certificate?	If the external CA configuration options - <code>ECA_CERT_PATH</code>, <code>ECA_PRIVATE_KEY_PATH</code>, and <code>ECA_TRUST_STORE_PATH</code> - are defined in the NetBackup configuration file of the source primary server, it can establish the trust using an external certificate. In the case of the Windows certificate trust store, only the option <code>ECA_CERT_PATH</code> is defined.
Which certificate authorities (CA) does the target primary server support?	The target primary server may support external CA, NetBackup CA, or both. See “View the Certificate authority for secure communication” on page 473.

The following table lists the CA support scenarios and the certificate to use to establish trust between the source and the target primary servers. The instructions assume that you use the NetBackup web UI for the configuration.

Table 32-2 Certificate to use for the trust setup

Can the primary server use an external certificate?	Which CA does the target primary server use?	Certificate to use for the trust setup
Yes	External CA	External CA
The source primary server can use the NetBackup CA and an external CA for communication with a remote primary server.	NetBackup CA	NetBackup CA
	External CA and NetBackup CA	NetBackup prompts to select the CA that you want to use for trust setup.
No	External CA	No trust is established.
The source primary server can only use the NetBackup CA for communication with a remote primary server.	NetBackup CA	NetBackup CA
	External CA and NetBackup CA	NetBackup CA

Add a trusted primary server

Replication operations require that a trust relationship exists between the NetBackup servers in the different domains. You can create a trust relationship between the primary servers that both use the NetBackup CA or that both use an external CA.

Before you begin, review the following information:

- Ensure that you have the RBAC System Administrator role or a role with similar permissions. Or, for appliances with software versions 3.1 and later you must have permissions for the NetBackup CLI user.
- For a remote Windows primary server, the user's domain may not be the same as that of the authentication service. In this case you must add the domain with LDAP using the `evssat addldapdomain` command.
- For a NetBackup CA-signed certificate, the recommended method to authenticate the server is the option **Specify authentication token of the trusted primary server**.
- If you use the option **Specify credentials of the trusted primary server**, that method may present a possible security breach. Only an authentication token can provide restricted access and allow secure communication between both the hosts. To establish trust with a 3.1 NetBackup primary appliance, use the NetBackup CLI credentials.

To add a trusted primary server

- 1 Open the NetBackup web UI.
- 2 Identify the NetBackup versions that are installed and the certificate types that are used on the source and the target servers.

The NetBackup web UI does not support adding a trusted primary that uses NetBackup version 8.0 or earlier. Both servers must use the same certificate type.
- 3 For the servers that use the NetBackup certificate authority (CA), obtain an authorization token for the remote server.

See [“Manage NetBackup certificate authorization tokens”](#) on page 427.
- 4 For the servers that use the NetBackup certificate authority (CA), obtain the fingerprint for each server.

See [“Manage NetBackup security certificates”](#) on page 424.
- 5 At the top right, select **Settings > Global security**.
- 6 Select the **Trusted primary servers** tab.
- 7 Select the **Add** button.
- 8 Enter the fully-qualified host name of the remote primary server and select **Validate Certificate Authority**.
- 9 Follow the prompts in the wizard.
- 10 Repeat these steps on the remote primary server.

More information

For more information on using an external CA with NetBackup, see the [NetBackup Security and Encryption Guide](#).

Remove a trusted primary server

Note: Any trusted primary servers at NetBackup version 8.0 or earlier must be removed using the NetBackup Administration Console or the NetBackup CLI.

You can remove a trusted primary server, which removes the trust relationship between primary servers. Note the following implications:

- Any replication operations fail that require the trust relationship.
- A remote primary server is not included in any usage reporting after you remove the trust relationship.

To remove a trusted primary server, you must perform the following procedure on both the source and the target server.

To remove a trusted primary server

- 1 Open the NetBackup web UI.
- 2 Ensure that all replication jobs to the target primary server are complete.
- 3 Delete all storage lifecycle policies (SLPs) that use the trusted primary as a destination. Before deleting an SLP, ensure that there are no backup policies or protection plans that use the SLP for storage.
- 4 At the top right, select **Settings > Global security**.
- 5 Select the **Trusted primary servers** tab.
- 6 Locate the server that you want to remove.
- 7 Select **Actions > Remove**.
- 8 Select **Remove trust**.

Note: If you use multiple NICs, if you established trust using more than one host NIC and if you remove the trust relationship with any one host NIC, the trust with all the other host NICs is broken.

Configure the audit retention period

Use the NetBackup web UI to set the retention period for the audit records, in days. The default retention period for the audit records is 90 days.

To configure audit retention period

- 1 On the top right, select **Settings > Global security**.
- 2 Go to the **Security controls** tab and locate the **Audit records retention period** section.
- 3 Enter the retention period. The value represents days should be either 0 (zero) or more than 27. The value 0 indicates that the audit records are never deleted.

Using access keys, API keys, and access codes

This chapter includes the following topics:

- [Access keys](#)
- [API keys](#)
- [Access codes](#)

Access keys

NetBackup access keys provide access the NetBackup interfaces through API keys and access codes.

See [“API keys”](#) on page 485.

See [“Access codes”](#) on page 491.

API keys

A NetBackup API key is a pre-authenticated token that identifies a NetBackup user to NetBackup RESTful APIs. The user can use the API key in an API request header when a NetBackup API requires authentication. API keys can be created for authenticated NetBackup users (groups are not supported). A specific API key is only created one time and cannot be recreated. Each API key has a unique key value and API key tag. NetBackup audits operations that are performed with that key with the full identity of the user.

The 'View' RBAC permission is required to create an API key.

The following actions are available for administrators and API key users.

- Administrators with the applicable role or RBAC permissions can manage API keys for all users. These roles are the Administrator, the Default Security Administrator, or a role with RBAC permissions for API keys.
- An authenticated NetBackup user can add and manage their own API key in the NetBackup web UI. If a user does not have access to the web UI, they can use the NetBackup APIs to add or manage a key.

Note: Starting with NetBackup 10.5, if multi-person authorization is enabled for API key operations, a ticket is generated. After the multi-person authorization ticket is approved, the user needs to execute the ticket using the **Execute ticket** option in the NetBackup web UI and then the required API key operation is executed.

For NetBackup releases earlier than 10.5, if multi-person authorization is enabled, you cannot perform API key operations.

More information

See [“User identity in the audit report”](#) on page 416.

See the [NetBackup Security and Encryption Guide](#) for information on using API keys with the `bpnbat` command.

Add an API key or view API key details (Administrators)

The API key administrator can manage the keys that are associated with all NetBackup users.

Add an API key

Note: Only one API key can be associated with a specific user at a time. If a user requires a new API key, the user or administrator must delete the key for that user. An expired API key can be reissued. The 'View' RBAC permission is required to create an API key.

To add an API key

- 1 On the left, select **Security > Access keys**. Then select the **API keys** tab.
- 2 Select the **Add** button.
- 3 Enter a **Username** for which you want to create the API key.
- 4 (Conditional) If the API key is for a SAML user, select **SAML authentication**.

A new API key for a SAML user remains inactive until the user signs into the web UI.

- 5 Indicate how long you want the API key to be valid, from today's date.

NetBackup calculates the expiration date and displays it.

- 6 Select the **Add** button.

- 7 To copy the API key, select **Copy and close**.

Store this key in a safe place. After you select **Copy and close**, the key cannot be retrieved again. If this API key replaces a previous key for your account, you must update any scripts, etc. to reflect the new API key.

View API key details

An API key administrator can view the API key details that are associated with all NetBackup users.

To view API key details

- 1 On the left, select **Security > Access keys**. Then select the **API keys** tab.
- 2 Locate the API key that you want to view.
- 3 To edit the date or description for the key, select the check box for the key. Then select **Actions > Edit**.

Edit, reissue, or delete an API key (Administrators)

As an API key administrator, you can edit API key details and reissue or delete API keys.

Note: Starting with NetBackup 10.5, if multi-person authorization is enabled for API key operations, a ticket is generated. After the multi-person authorization ticket is approved, editing, reissuing, or deleting an API key is performed. For NetBackup releases earlier than 10.5, if multi-person authorization is enabled, you cannot perform API key operations.

Edit the expiration date or description for an API key

Note: For SAML users, avoid selecting an expiration date for the API key that occurs after the SAML session expires. If the date occurs after the session expires, this action can introduce a security risk with that API key.

You can edit the description of an API key or change the expiration date of an active API key.

To edit the expiration date or description for an API key

- 1 On the left, select **Security > Access keys > API keys**.
- 2 Select the check box for the API key that you want to edit.
- 3 Select **Actions > Edit**.
- 4 Note the current expiration date for the key and extend the date as wanted.
- 5 Make any wanted changes to the description.
- 6 Select **Save**.

Reissue an API key after it expires

Note: For SAML users, avoid selecting an expiration date for the API key that occurs after the SAML session expires. If the date occurs after the session expires, this action can introduce a security risk with that API key.

When an API key expires you can reissue the API key. This action creates a new API key for the user.

To reissue an API key

- 1 On the left, select **Security > Access keys > API keys**.
- 2 Select the check box for the API key that you want to edit.
- 3 Select the **Actions** menu. Then select **Reissue > Reissue**.

Delete an API key

You can delete an API key to remove access for the user or when the key is no longer used. The key is permanently deleted, meaning that the associated user can no longer use that key for authentication.

To delete an API key

- 1 On the left, select **Security > Access keys > API keys**.
- 2 Locate and select the API key that you want to delete.
- 3 Select the **Actions** menu. Then select **Delete > Delete**.

Add an API key or view your API key details

You can create an API key to authenticate your NetBackup user account when using NetBackup RESTful APIs.

Add an API key

As a NetBackup web UI user you can use the web UI to add or view the details for your own API key.

To add an API key

- 1 If your API key has expired you can reissue the key.
See [the section called “Reissue your API key after it expires”](#) on page 490.
- 2 On the top right, select the profile icon and select **Add API key**.
- 3 (Non-SAML users) Indicate how long you want the API key to be valid, from today's date.
NetBackup calculates the expiration date and displays it.
- 4 (SAML users) After NetBackup validates the token from the SAML session, then the expiration date for the API key can be determined.
- 5 Select the **Add** button.
- 6 To copy the API key, select **Copy and close**.
Store this key in a safe place. After you select **Copy and close**, the key cannot be retrieved again. If this API key replaces a previous key for your account, you must update any scripts, etc. to reflect the new API key.

View your API key details

To view your API key details

- ◆ On the top right, select the profile icon and select **View my API key details**.

Edit, reissue, or delete your API key

You can manage your own API key from the NetBackup web UI.

Edit the expiration date or description for your API key (non-SAML users)

Non-SAML users can change the expiration date for an active API key. After an API key expires, you can reissue the key.

To edit your API key details

- 1 On the top right, click the profile icon and click **View my API key details**.
Note: If your API key is expired, you can click **Reissue** to reissue the key.
See [the section called “Reissue your API key after it expires”](#) on page 490.
- 2 Click **Edit**.
- 3 Note the current expiration date for the key and extend the date as wanted.
- 4 Make any wanted changes to the description.
- 5 Click **Save**.

Reissue your API key after it expires

When your API key expires you can reissue the API key. This action creates a new API key for you.

To reissue your API key

- 1 On the top right, click the profile icon and click **View my API key details**.
- 2 On the top right, click **Reissue**.
- 3 (Non-SAML users) Note the current expiration date for the key and extend the date as wanted.
- 4 Make any wanted changes to the description.
- 5 Click **Reissue**.

Delete your API key

You can delete an API key if you no longer have access to the key or no longer use it. When you delete an API key, that key is permanently deleted. You can no longer use that key for authentication or with the NetBackup APIs.

To delete your API key

- 1 On the top right, click the profile icon and click **View my API key details**.
- 2 On the top right, click **Delete**. Then click **Delete**.

Use an API key with NetBackup REST APIs

After a key is created, the user can pass the API key in the API request headers. For example:

```
curl -X GET https://primaryservername.domain.com/netbackup/admin/jobs/5 \
-H 'Accept: application/vnd.netbackup+json;version=3.0' \
-H 'Authorization: <API key value>'
```

Access codes

To run certain NetBackup administrator commands, for example `bpererror`, you need to authenticate through the web UI. You need to generate an access code through the command-line interface, get the access request approved from the administrator, and then access the command.

With the web UI authentication for CLI access, NetBackup administrators can delegate the associated privileges to other users. By default, only a root administrator or an administrator can perform NetBackup operations through the command-line interface. The web UI authentication support allows non-root users to administer NetBackup who have CLI access that the Security Administrator has granted. You can also administer NetBackup with a non-RBAC user role (such as Operating System Administrator) even though you are not registered as a NetBackup user. Each time you need to generate a new access code to access CLIs.

Request CLI access through web UI authentication

To run NetBackup commands using the NetBackup CLI, the following requirements exist for the user:

- The user must have the RBAC role Default NetBackup Command Line (CLI) Administrator or a role with similar permissions.
- The user must submit a request for temporary access to the CLI. By default, a CLI access session is valid for 24 hours.

The command that the user runs for the request depends on whether or not they have access to the NetBackup web UI.

See [the section called “Request CLI access when you have access to the NetBackup web UI”](#) on page 491.

See [the section called “Request CLI access from the security administrator”](#) on page 492.

Request CLI access when you have access to the NetBackup web UI

If you have access to the NetBackup web UI, you can use the web UI to approve a CLI access request using the access code from the `bpnbat` command.

To request CLI access

- 1 Run the following command:

```
bpnbat -login -logintype webui
```

An access code is generated.
- 2 Open the NetBackup web UI.
- 3 On the top right, select the profile icon.
- 4 Select **Approve access request**.
- 5 Enter the CLI access code that was created when you ran the `bpnbat` command. Then select **Review**.
- 6 Review the access request details.
- 7 Select **Approve**.
- 8 After you approve the request, you can use the command-line interface to run the wanted commands.

Request CLI access from the security administrator

If you do not have access to the NetBackup web UI, you must submit a request for a CLI access to the security administrator. A user with the Default Security Administrator role or a role with similar permissions must approve the request.

To request CLI access from the security administrator

- 1 Run the following command:

```
bpnbat -login -logintype webui -requestApproval
```

An access code is generated.
- 2 Contact the security administrator and give them the access code to approve the CLI access request.

See [“Approve the CLI access request of another user”](#) on page 492.
- 3 After the request is approved, you can use the command-line interface to run the wanted commands.

Approve the CLI access request of another user

If you have the Default Security Administrator role or a role with similar permissions, you can approve the request of another user who needs CLI access. Note that to run commands, that user must also have the RBAC role Default NetBackup Command Line (CLI) Administrator or a role with similar permissions.

To approve the CLI access request of another user

- 1 The user that requires CLI access must first run the following command to request approval:

```
bpnbat -login -logintype webui -requestApproval
```
- 2 Sign in to the NetBackup web UI.
- 3 On the left, select **Security > Access keys**. Then select the **Access codes** tab.
- 4 Enter the CLI access code that you have received from the user who requires CLI access and select **Review**.
- 5 Review the access request details.
- 6 (Optional) Provide any comments.
- 7 Select **Approve**.

Edit the settings for command-line access

You can configure the default time that is set for a CLI session when a user requests CLI access.

To edit the settings for command-line access

- 1 On the left, select **Security > Access keys**.
- 2 On the right, select **Access settings**.
- 3 Select **Edit**.
- 4 Enter the time in minutes or hours that you want the CLI access session to be valid. 1 minute is the minimum value and 24 hours is the maximum value.

Configuring authentication options

This chapter includes the following topics:

- [Sign-in options for the NetBackup web UI](#)
- [Configure user authentication with smart cards or digital certificates](#)
- [About single sign-on \(SSO\) configuration](#)
- [Configure NetBackup for single sign-on \(SSO\)](#)
- [Troubleshooting SSO](#)

Sign-in options for the NetBackup web UI

NetBackup supports authentication of local domain users and Active Directory (AD) or LDAP domain users. AD and LDAP domains, smart card, and single sign-on (SSO with SAML) requires separate configuration for each primary server domain where you want to use the authentication method.

NetBackup supports the following types of user authentication:

- Username and password
- Digital certificate or smart card, including CAC and PIV
This authentication method only supports one AD or LDAP domain for each primary server domain and is not available for local domain users.
See [“Configure user authentication with smart cards or digital certificates”](#) on page 495.
- Single sign-on, with SAML
Note the following requirements and limitations.

- To use SSO, you must have a SAML 2.0 compliant identity provider configured in your environment.
- Only one AD or LDAP domain is supported for each primary server domain. This feature is not available for local domain users.
- Configuration of the IDP requires the NetBackup APIs or the NetBackup command `nbidpcmd`.
- API keys are used to authenticate a user or a group and cannot be used with SAML-authenticated users or groups.
- Global logout is not supported.

See [“Configure NetBackup for single sign-on \(SSO\)”](#) on page 501.

Configure user authentication with smart cards or digital certificates

You can map a smart card or certificate with an AD or an LDAP domain for user validation. Alternatively, you can configure a smart card or certificate without an AD or an LDAP domain.

See [“Configure smart card authentication with a domain”](#) on page 495.

See [“Configure smart card authentication without a domain”](#) on page 496.

Configure smart card authentication with a domain

You can configure NetBackup to validate users with smart cards or certificates with an AD or an LDAP domain.

Note the following prerequisites:

- Before you add the authentication method you must add the domain that is associated with your NetBackup users. See the [NetBackup Security & Encryption Guide](#).
- Ensure that you complete the role-based access control (RBAC) configuration for the NetBackup users before you configure smart card or certificate authentication.

See [“Configuring RBAC”](#) on page 517.

To configure smart card authentication with a domain

- 1 Sign in to the NetBackup web UI.
- 2 At the top right, select **Settings > Smart card authentication**.
- 3 Turn on **Smart card authentication**.

- 4** Select the required AD or LDAP domain from the **Select the domain** option.
- 5** Select a **Certificate mapping attribute**: Common name (CN) or Universal principal name (UPN).
- 6** Optionally, enter the **OCSP URI**.
If you do not provide the OCSP URI, the URI in the user certificate is used.
- 7** Select **Save**.
- 8** To the right of **CA certificates**, click **Add**.
- 9** Browse for or drag and drop the **CA certificates** and click **Add**.

Smart card authentication requires a list of trusted root or intermediate CA certificates. Add the CA certificates that are associated with the user digital certificates or the user smart cards.

Certificate file types must be .crt, .cer, .der, .pem, or PKCS #7 format and less than 64KB in size.
- 10** On the **Smart card authentication** page, verify the configuration information.
After configuring smart card authentication, you must restart the NetBackup Web Management Console (nbwmc) service.
- 11** Before users can use a digital certificate that is not installed on a smart card, the certificate must be uploaded to the browser's certificate manager.

See the browser documentation for instructions or contact your certificate administrator for more information.
- 12** When users sign in, they now see an option to **Sign in with certificate or smart card**.

If you do not want users to have this sign-in option yet, turn off **Smart card authentication**. (For example, if all users do not yet have their certificates configured on their hosts.). The settings that you configured are retained even if you turn off smart card authentication.

For such users, the domain name and domain type are smart card.

Configure smart card authentication without a domain

You can configure NetBackup to validate users with smart cards or certificates without an associated AD or LDAP domain. Only users are supported for this configuration. User groups are not supported.

To configure smart card authentication without a domain

- 1 At the top right, select **Settings > Smart card authentication**.
- 2 Turn on **Smart card authentication**.
- 3 (Conditional step) If AD or LDAP domain is configured in your environment, select **Continue without the domain** option.
- 4 Select a **Certificate mapping attribute**: Common name (CN) or Universal principal name (UPN).
- 5 Optionally, enter the **OCSP URI**.
 If you do not provide the OCSP URI, the URI in the user certificate is used.
- 6 Select **Save**.
- 7 To the right of **CA certificates**, click **Add**.
- 8 Browse for or drag and drop the **CA certificates** and click **Add**.
- 9 Smart card authentication requires a list of trusted root or intermediate CA certificates. Add the CA certificates that are associated with the user digital certificates or the user smart cards.
 Certificate file types must be `.crt`, `.cer`, `.der`, `.pem`, or PKCS #7 format and less than 64KB in size.
- 10 On the **Smart card authentication** page, verify the configuration information.
 After configuring smart card authentication, you must restart the NetBackup Web Management Console (nbwmc) service.
 Before users can use a digital certificate that is not installed on a smart card, the certificate must be uploaded to the browser's certificate manager.
- 11 When users sign in, they now see an option to **Sign in with certificate or smart card**.
 If you do not want users to have this sign-in option yet, turn off **Smart card authentication**. (For example, if all users do not yet have their certificates configured on their hosts.). The settings that you configured are retained even if you turn off smart card authentication.

Edit the configuration for smart card authentication

If the configuration changes for smart card authentication, you can edit the configuration details.

To edit user authentication configuration with domain

- 1 At the top right, select **Settings > Smart card authentication**.
- 2 You may want to edit the AD or LDAP domain selection in the following cases:
 - To select a domain that is different than the existing one
 - The existing domain is deleted and you want to select a new domain
 - You want to continue without the domain
 Select **Edit**.
- 3 Select a domain.
 Only the domains that are configured for NetBackup display in this list.
 If you do not want to validate the users with domain, you can select **Continue without the domain**.
- 4 Edit the **Certificate mapping attribute**.
- 5 Leave the **OCSP URI** field empty if you want to use the **URI** value from the user certificate. Or, provide the URI that you want to use.

Add or delete a CA certificate that is used for smart card authentication

Add a CA certificate

Smart card authentication requires a list of trusted root or intermediate CA certificates. Add the CA certificates that are associated with the user digital certificates or the user smart cards.

To add a CA certificate

- 1 At the top right, select **Settings > Smart card authentication**.
- 2 Locate **CA certificates** and select the **Add** button.
- 3 Browse for or drag and drop the **CA certificates**. Then select the **Add** button.

Smart card authentication requires a list of trusted root or intermediate CA certificates. Add the CA certificates that are associated with the user digital certificates or the user smart cards.

Certificate file types must be in DER, PEM, or PKCS #7 format and no more than 1 MB in size.

Delete a CA certificate

You can delete a CA certificate if it is no longer used for smart card authentication. Note that if a user attempts to use the associated digital certificate or smart card certificate, they are not able to sign in to NetBackup.

To delete a CA certificate

- 1 At the top right, select **Settings > Smart card authentication**.
- 2 Select the CA certificates that you want to delete.
- 3 Select **Delete > Delete**.

Disable or temporarily disable smart card authentication

You can disable smart card authentication if you no longer want to use that authentication method for the primary server. Or, if you need to complete other configuration before users can use smart cards.

To disable smart card authentication

- 1 At the top right, select **Settings > Smart card authentication**.
- 2 Turn off **Smart card authentication**.

The settings that you configured are retained even if you turn off smart card authentication.

About single sign-on (SSO) configuration

You can configure single sign-on (SSO) with any identity provider (IDP) that uses the SAML 2.0 protocol for exchanging authentication and authorization information. Note that you can configure an IDP with more than one Cohesity product. For example, the same IDP can be configured with NetBackup and with APTARE.

Note the following requirements and limitations:

- To use SSO, you must have a SAML 2.0 compliant identity provider configured in your environment.
- Only identity providers that use AD or LDAP directory services are supported.
- Configuration of the IDP requires the NetBackup APIs or the NetBackup command `nbidpcmd`.
- SAML users cannot use the APIs. API keys are used to authenticate a user and therefore cannot be used with a SAML-authenticated user.
- Global logout is not supported.

Figure 34-1 Example NAT configuration: Identity provider in a private network




Figure 34-2 Example NAT configuration: Primary server in a private network




Figure 34-3 Example configuration: Primary server and identity provider in same network





Figure 34-4 Example configuration: Primary server in private network and identity provider in public network



Configure NetBackup for single sign-on (SSO)

This section provides steps to set up trust and exchange configuration information between the IDP and the NetBackup primary server. Before proceeding with the steps, ensure that the following prerequisites are met in your environment:

- An IDP is set up and deployed in your environment.
- The IDP is configured to authenticate domain users of Active Directory (AD) or Lightweight Directory Access Protocol (LDAP).

Table 34-1 Steps to configure NetBackup for single sign-on

Step	Action	Description
1.	Download the IDP metadata XML file	Download and save the IDP metadata XML file from the IDP. SAML metadata that is stored in XML files is used to share configuration information between the IDP and the NetBackup primary server. The IDP metadata XML file is used to add the IDP configuration to the NetBackup primary server.
2.	Configure the SAML keystore, and add and enable the IDP configuration on the NetBackup primary server	See “Configure the SAML KeyStore” on page 502. See “Configure the SAML keystore and add and enable the IDP configuration” on page 505.

Table 34-1 Steps to configure NetBackup for single sign-on (*continued*)

Step	Action	Description
3.	Download the service provider (SP) metadata XML file	The NetBackup primary server is the SP in the NetBackup environment. You can access the SP metadata XML file from the NetBackup primary server by entering the following URL in your browser: <code>https://primaryserver/netbackup/sso/saml2/metadata</code> Where <i>primaryserver</i> is the IP address or host name of the NetBackup primary server.
4.	Enroll the NetBackup primary server as a service provider (SP) with the IDP	See “Enroll the NetBackup primary server with the IDP” on page 508.
5.	Add SAML users and the SAML groups that use SSO to the necessary RBAC roles	SAML users and SAML user groups are available in RBAC only if the IDP is configured and enabled on the NetBackup primary server. For steps on adding RBAC roles, see the following topic. See “Add a user to a role (non-SAML)” on page 519.

After the initial setup, you can choose to enable, update, disable, or delete the IDP configuration.

See [“Manage an IDP configuration”](#) on page 509.

After the initial setup, you can choose to update, renew, or delete the NetBackup CA SAML keystore . You can also configure and manage the ECA SAML keystore.

Configure the SAML KeyStore

To establish a trust between the NetBackup primary server and the IDP server, you must configure an SAML KeyStore on the NetBackup primary server. Depending on whether you are using the NetBackup CA or an external certificate authority (ECA), refer to either of the following sections:

Note: If you are using a combination of an ECA and NetBackup CA in your environment, by default, the ECA is considered while establishing trust with the IDP server.

Note: The SAML KeyStore configuration using batch files, such as `configureCerts.bat`, `configureCerts`, `configureSAMLECACert.bat`, `configureSAMLECACert` and their corresponding options is deprecated.

Configure a NetBackup CA KeyStore

If you are using the NetBackup CA, create the NetBackup CA KeyStore on the NetBackup primary server.

To create a NetBackup CA KeyStore

- 1 Log on to the NetBackup primary server as root or administrator.
- 2 Run the following command:

```
nbidpcmd -cCert -M master_server -f
```

`-f` is optional. Use the option for the forceful update.

Once the NetBackup CA KeyStore is created, ensure that you update the NetBackup CA KeyStore every time the NetBackup CA certificate is renewed.

To renew the NetBackup CA KeyStore

- 1 Log on to the NetBackup primary server as root or administrator.
- 2 Run the following command:

```
nbidpcmd -rCert -M master_server
```

- 3 Download the new SP metadata XML file from the NetBackup primary server by entering the following URL in your browser:

`https://primaryserver/netbackup/sso/saml2/metadata`

Where *primaryserver* is the IP address or host name of the NetBackup primary server.

- 4 Upload the new SP metadata XML file to the IDP.

See [“Enroll the NetBackup primary server with the IDP”](#) on page 508.

To remove the NetBackup CA KeyStore

- 1 Log on to the NetBackup primary server as root or administrator.
- 2 Run the following command

```
nbidpcmd -dCert -M master_server
```

- 3 Download the new SP metadata XML file from the NetBackup primary server by entering the following URL in your browser:

`https://primaryserver/netbackup/sso/saml2/metadata`

Where *primaryserver* is the IP address or host name of the NetBackup primary server.

- 4 Upload the new SP metadata XML file to the IDP.
- 5 See [“Enroll the NetBackup primary server with the IDP”](#) on page 508.

Configure an ECA KeyStore

If you are using an ECA, import the ECA KeyStore to the NetBackup primary server.

Note: If you are using a combination of an ECA and the NetBackup CA in your environment, by default, the ECA is considered while establishing trust with the IDP server. To use the NetBackup CA, you must first remove the ECA KeyStore.

To configure an ECA KeyStore

- 1 Log on to the primary server as root or administrator.
- 2 Depending on whether you want to configure SAML ECA keystore using the configured NetBackup ECA KeyStore or you want to provide the ECA certificate chain and private key, run the following commands:
 - Run the following command to use NetBackup ECA configured KeyStore:
`nbidpcmd -cECACert -uECA existing ECA configuration [-f] [-M primary_server]`
 - Run the following command to use ECA certificate chain and private key provided by the user:
`nbidpcmd -cECACert -certPEM certificate chain file -privKeyPath private key file [-ksPassPath Keystore Passkey File] [-f] [-M <master_server>]`
 - Certificate chain file specifies the certificate chain file path. The file must be in PEM format and must be accessible to the primary server on which the configuration is being performed.
 - Private key file specifies the private key file path. The file must be in PEM format and must be accessible to the primary server on which the configuration is being performed.
 - KeyStore passkey file specifies the KeyStore password file path and must be accessible to the primary server on which the configuration is being performed.

- Primary server is the host name or IP address of primary server on which you want to perform SAML ECA KeyStore configuration. The NetBackup primary server where you run the command is selected by default.

To remove the ECA KeyStore

- 1 Log on to the primary server as root or administrator.
- 2 Download the new SP metadata XML file from the NetBackup primary server by entering the following URL in your browser:

`https://primaryserver/netbackup/sso/saml2/metadata`

Where *primaryserver* is the IP address or host name of the NetBackup primary server.

- 3 Upload the new SP metadata XML file to the IDP.

See [“Enroll the NetBackup primary server with the IDP”](#) on page 508.

Configure the SAML keystore and add and enable the IDP configuration

Before proceeding with the following steps, ensure that you have downloaded the IDP metadata XML file and saved it on the NetBackup primary server.

To configure SAML keystore and add and enable an IDP configuration

- 1 Log on to the primary server as root or administrator.
- 2 Run the following command.

For IDP and NetBackup CA SAML KeyStore configuration:

```
nbidpcmd -ac -n IDP configuration name -mxp IDP XML metadata file
[-t SAML2] [-e true | false] [-u IDP user field] [-g IDP user
group field] [-cCert] [-f] [-M primary server]
```

Alternatively for IDP and ECA SAML KeyStore configuration:

Depending on whether you want to configure SAML ECA KeyStore using the configured NetBackup ECA KeyStore or you want to provide the ECA certificate chain and private key, run the following commands:

- Use NetBackup ECA configured keystore:

```
nbidpcmd -ac -n IDP configuration name -mxp IDP XML metadata
file [-t SAML2] [-e true | false] [-u IDP user field] [-g IDP
user group field] -cECACert -uECA existing ECA configuration
[-f] [-M Primary Server]
```

- Use ECA certificate chain and private key provided by the user:

```
nbidpcmd -ac -n IDP configuration name -mxp IDP XML metadata file [-t SAML2] [-e true | false] [-u IDP user field] [-g IDP user group field] -cECACert -certPEM certificate chain file -privKeyPath private key file [-ksPassPath KeyStore passkey file] [-f] [-M primary server]
```

Replace the variables as follows:

- *IDP configuration name* is a unique name provided to the IDP configuration.
- *IDP XML metadata file* is the path to the XML metadata file, which contains the configuration details of the IDP in Base64URL-encoded format.
- *-e true | false* enables or disables the IDP configuration. An IDP configuration must be added and enabled, otherwise users cannot sign in with the single sign-on (SSO) option. Even though you can add multiple IDP configurations on a NetBackup primary server, only one IDP configuration can be enabled at a time.
- The SAML attribute names *IDP user field* and *IDP user group field* are used to map user identity information and group information in the Identity Provider. These fields are optional, and if not provided, they are mapped to the `userPrincipalName` and `memberOf` SAML attributes by default. For instance, if you have customized the attribute mapping in the Identity Provider to use attributes like email and groups, when configuring the SAML configuration, you need to provide the *-u* option for email and *-g* option for groups.

If you have not provided values for these attributes during configuration, ensure that the Identity Provider returns the values against the `userPrincipalName` and `memberOf` attributes.

For Example:

If SAML response is as follows:

```
saml:AttributeStatement <saml:Attribute Name="userPrincipalName">
<saml:AttributeValue>username@domainname</saml:AttributeValue>
</saml:Attribute> <saml:Attribute Name="memberOf">
<saml:AttributeValue>CN=group name,
DC=domainname</saml:AttributeValue> </saml:Attribute>
</saml:AttributeStatement>
```

It implies that you need to map the *-u* and *-g* options against the fields "saml:Attribute Name".

Note: Ensure that the SAML attribute values are returned in the format of *username@domainname* for the field mapped to the `-u` option that defaults to `userPrincipalName`. If you include the domain name when returning group information, it should follow the format "(CN=group name, DC=domainname)" or "(domainname\groupname)".

However, if you return the group name as plain text without domain information, it should be mapped without the domain name in the SAML RBAC group.

- *primary Server* is the host name or IP address of primary server to which you want to add or modify the IDP configuration. The NetBackup primary server where you run the command is selected by default.
- *Certificate Chain File* is the certificate chain file path. The file must be in PEM format and must be accessible to the primary server on which the configuration is being performed.
Private Key File is the private key file path. The file must be in PEM format and must be accessible to the primary server on which the configuration is being performed.
KeyStore Passkey File is the KeyStore passkey file path and must be accessible to the primary server on which the configuration is being performed.

If your Identity Provider is already configured with SAML attribute names as `userPrincipalName` and `memberOf`, you do not have to provide the `-u` and `-g` option while configuration. If you are using any other custom attributes name, provide those names against `-u` and `-g` as follows:

For example:

If the Identity Provider SAML attribute names are mapped as "email" and "groups", use the following command for configuration:

```
nbidpcmd -ac -n veritas_configuration -mxp file.xml -t SAML2 -e true -u email -g groups -cCert -Mprimary_server.abc.com
```

`-u` and `-g` are optional and it depends on the Identity Provider configuration. Ensure that you specify the same parameter values that you have provided at the time of configuration.

Enroll the NetBackup primary server with the IDP

The NetBackup primary server must be enrolled with the IDP as a service provider (SP). For step-by-step procedures that are specific to a particular IDP, see the following table:

Table 34-2 IDP-specific steps for enrolling the NetBackup primary server

IDP name	Link to steps
ADFS	https://www.veritas.com/docs/100047744
Okta	https://www.veritas.com/docs/100047745
PingFederate	https://www.veritas.com/docs/100047746
Azure	https://www.veritas.com/docs/100047748
Shibboleth	https://www.veritas.com/docs/00047747

Enrolling an SP with an IDP typically involves the following operations:

Uploading the SP metadata XML file to the IDP

The SP metadata XML file contains the SP certificate, the entity ID, the Assertion Consumer Service URL (ACS URL), and a log out URL (SingleLogoutService). The SP metadata XML file is required by the IDP to establish trust, and exchange authentication and authorization information with the SP.

Mapping the SAML attributes to their AD or LDAP attributes

Attribute mappings are used to map SAML attributes in the SSO with its corresponding attributes in the AD or LDAP directory. The SAML attribute mappings are used for generating SAML responses, which are sent to the NetBackup primary server. Ensure that you define SAML attributes that map to the `userPrincipalName` and the `memberOf` attributes in the AD or LDAP directory. The SAML attributes must adhere to the following formats:

Table 34-3

Corresponding AD or LDAP attribute	SAML attribute format
<code>userPrincipalName</code>	<code>username@domainname</code>
<code>memberOf</code>	<code>(CN=group name, DC=domainname)</code>

Note: While adding the IDP configuration to the NetBackup primary server, the values entered for the user (`-u`) and user group (`-g`) options must match the SAML attribute names that are mapped to the `userPrincipalName` and the `memberOf` attributes in the AD or LDAP.

See [“Configure the SAML keystore and add and enable the IDP configuration”](#) on page 505.

Manage an IDP configuration

You can manage the identity provider (IDP) configurations on the NetBackup primary server by using the enable (`-e true`), update (`-uc`), disable (`-e false`), and delete (`-dc`) options of the `nbidpcmd` command.

Enable an IDP configuration

By default, an IDP configuration is not enabled in the product environment. If you did not enable the IDP when you added it, you can use the `-uc -e true` options to update and enable the IDP configuration.

To enable an IDP configuration

- 1 Log on to the primary server as root or administrator.
- 2 Run the following command:

```
nbidpcmd -uc -n IDP configuration name -e true
```

Where *IDP configuration name* is a unique name provided to the IDP configuration.

Note: Even though you can configure multiple IDPs on a NetBackup primary server, only one IDP can be enabled at a time.

Update an IDP configuration

You can update the XML metadata file associated with an IDP configuration.

To update the IDP XML metadata file in an IDP configuration

- 1 Log on to the primary server as root or administrator.
- 2 Run the following command:

```
nbidpcmd -uc -n IDP configuration name -mxp IDP XML metadata file
```

Replace the variables as described below:

- *IDP configuration name* is a unique name provided to the IDP configuration.

- *IDP XML metadata file* is the path to the XML metadata file, which contains the configuration details of the IDP in Base64URL-encoded format.

If you want to update the IDP user or IDP user group values in an IDP configuration, you must first delete the configuration. The single sign-on (SSO) option is not available for users until you re-add the configuration with the updated IDP user or IDP user group values.

To update IDP user or IDP user group in an IDP configuration

- 1 Log on to the primary server as root or administrator.
- 2 Delete the IDP configuration.

```
nbidpcmd -dc -n IDP configuration name
```

Where *IDP configuration name* is a unique name provided to the IDP configuration.

- 3 To add and enable the configuration again, run the following command:

```
nbidpcmd -ac -n IDP configuration name -mxp IDP XML metadata file  
[-t SAML2] [-e true | false] [-u IDP user] [-g IDP user group  
field] [-M Master Server]
```

Replace the variables as described below:

- *IDP configuration name* is a unique name provided to the IDP configuration.
- *IDP XML metadata file* is the path to the XML metadata file, which contains the configuration details of the IDP in Base64URL-encoded format.
- `-e true | false` enables or disables the IDP configuration. An IDP must be available and enabled otherwise users cannot sign in with the single sign-on (SSO) option. Even though you can add multiple IDP configurations on a NetBackup primary server, only one IDP configuration can be enabled at a time.
- *Master Server* is the host name or IP address of the primary server to which you want to add or modify the IDP configuration. The NetBackup primary server where you run the command is selected by default.

Disable an IDP configuration

If an IDP configuration is disabled in the product environment, the single sign-on (SSO) option of that IDP is not available for users when they sign in.

To disable an IDP configuration

- 1 Log on to the primary server as root or administrator.
- 2 Run the following command:

```
nbidpcmd -uc -n IDP configuration name -e false
```

Where *IDP configuration name* is a unique name provided to the IDP configuration.

Delete an IDP configuration

If an IDP configuration is deleted, the single sign-on (SSO) option of that IDP is not available for users when they sign in.

To delete an IDP configuration

- 1 Log on to the primary server as root or administrator.
- 2 Run the following command:

```
nbidpcmd -dc -n IDP configuration name
```

Where *IDP configuration name* is a unique name provided to the IDP configuration.

Video: Configure single sign-on in NetBackup

In this video, you will see an overview of how to configure single sign-on (SSO) in NetBackup.

Video link

Depending on which IDP you are using, see the following articles for steps on downloading the IDP metadata XML file and enrolling the NetBackup primary server with the IDP:

- ADFS: <https://www.veritas.com/docs/100047744>
- Okta: <https://www.veritas.com/docs/100047745>
- PingFederate: <https://www.veritas.com/docs/100047746>
- Azure: <https://www.veritas.com/docs/100047748>
- Shibboleth: <https://www.veritas.com/docs/100047747>

More information is available about SSO in NetBackup.

See “Configure NetBackup for single sign-on (SSO)” on page 501.

Troubleshooting SSO

This section provides steps for troubleshooting issues related to SSO.

Redirection issues

If you are facing issues with redirection, check the error messages in web services log files to narrow down the cause of the issue. NetBackup creates logs for the NetBackup web server and for the web server applications. These logs are written to the following location:

- UNIX: `/usr/openv/logs/nbweb service`
- Windows:`install_path\NetBackup\logs\nbweb service`

NetBackup web UI does not redirect to the IDP sign in page

The IDP metadata XML file contains the IDP certificate, the entity ID, the redirect URL, and the logout URL. The NetBackup web UI can fail to redirect to the IDP sign in page, if the IDP XML metadata file is outdated or corrupted. The following message is added to the web service log:

`Failed to redirect to the IDP server.`

To ensure that the latest configuration details are available to the NetBackup primary server, download the latest copy of the XML metadata file from the IDP. Use the IDP XML metadata file to add and enable the latest IDP configuration on the NetBackup primary server. See [“Configure the SAML keystore and add and enable the IDP configuration”](#) on page 505.

IDP sign in page does not redirect to the NetBackup web UI

When you enter your credentials in the IDP sign in page, your browser might display an **Authentication failed** error, instead of redirecting to the NetBackup web UI. Refer to the following table for resolution steps based on the error found in the web service log.

Table 34-4

Web Service log error message	Explanation and recommended action
<code>userPrincipalName not found in response.</code>	While adding the IDP configuration to the NetBackup primary server, the value entered for the user (<code>-u</code>) option must match the SAML attribute name, which is mapped to the <code>userPrincipalName</code> attribute in AD or LDAP. For more information, See “Configure the SAML keystore and add and enable the IDP configuration” on page 505.

Table 34-4 (continued)

Web Service log error message	Explanation and recommended action
userPrincipalName is not in expected format	The IDP sends SAML responses to the NetBackup primary server, which contains SAML user and SAML user group information. To enable the IDP to successfully send this information, ensure the value of <code>userPrincipalName</code> attribute sent by the IDP is defined in the format of <code>username@domainname</code>. For more information, See “Enroll the NetBackup primary server with the IDP” on page 508.
Authentication issue instant is too old or in the future	This error can occur because of the following reasons: ■ The date and time of IDP server and the NetBackup primary server is not synchronized. ■ By default, the NetBackup primary server allows a user to remain authenticated for a period of 24 hours. You might encounter this error, If an IDP allows a user to remain authenticated for a period longer than 24 hours. To resolve this error, you can update the SAML authentication lifetime of the NetBackup primary server to match that of the IDP. Specify the new SAML authentication lifetime in the <code><installpath>\var\global\wsl\config\web.conf</code> file on the NetBackup primary server. For example, If your IDP has an authentication lifetime as 36 hours, update the entry in the <code>web.conf</code> file as follows: <pre>SAML_ASSERTION_LIFETIME_IN_SECS=129600</pre>
Response is not success	This error can occur because of the following reasons: ■ The IDP metadata XML file contains an IDP certificate. If you are using a NetBackup CA, ensure that the IDP certificate is updated with latest NetBackup CA certificate information. For more information, See “Configure the SAML KeyStore” on page 502. ■ The Certificate Revocation List (CRL) must be disabled in the IDP if you are using a NetBackup CA keystore.

Unable to sign in due to authorization-related issues

To sign in with SSO, you must add SAML users and the SAML user groups to the necessary RBAC roles. If the RBAC roles are not correctly assigned, you might encounter the following error while signing into NetBackup web UI.

You are not authorized to access this application. Contact your NetBackup security administrator to request RBAC permissions for the NetBackup web user interface.

Refer to the table below to troubleshoot authorization-related issues:

Table 34-5

Cause	Explanation and recommended action
RBAC roles are not assigned to the SAML users and the SAML groups.	After an IDP configuration is added and enabled on the NetBackup primary server, ensure that necessary RBAC roles are assigned to SAML users and SAML user groups that use SSO. Note that SAML users and SAML user groups are available in RBAC only after the IDP configuration is added and enabled on the NetBackup primary server. For steps on adding users, See “Add a user to a role (non-SAML)” on page 519.
RBAC roles are assigned to SAML users and SAML user groups associated with an IDP configuration that is not currently added and enabled.	When you add a SAML users or SAML user group in RBAC, the SAML user or SAML user group entry is associated with the IDP configuration that is added and enabled at that time. If you add and enable a new IDP configuration, ensure that you also add another entry for the SAML user or SAML user group. The new entry is associated with the new IDP configuration. For example, NBU_user is added to RBAC and assigned the necessary permissions, while an ADFS IDP configuration is added and enabled. If you add and enable an Okta IDP configuration, you must add a new user entry for NBU_user. Assign the necessary RBAC roles to the new user entry, which is associated with the Okta IDP configuration. For steps on adding users, See “Add a user to a role (non-SAML)” on page 519.
RBAC roles are assigned to local domain users or Active Directory (AD) or LDAP domain users (instead of SAML users and SAML user groups).	SAML user or SAML user group records might appear similar to corresponding local domain users or AD or LDAP domain users already added in the RBAC. After an IDP configuration is added and enabled on the NetBackup primary server, ensure that you add SAML users and SAML user groups in RBAC and assign the necessary permissions. Note that SAML users and SAML user groups are available in RBAC only after the IDP configuration is added and enabled on the NetBackup primary server. For steps on adding SAML users and user groups, See “Add a user to a role (non-SAML)” on page 519.

Table 34-5 (continued)

Cause	Explanation and recommended action
The NetBackup primary server is unable to retrieve user group information from the IDP	The IDP sends SAML responses to the NetBackup primary server, which contains SAML user and SAML user group information. To enable the IDP to successfully send this information, ensure the following: ■ The IDP is configured to authenticate domain users from AD or LDAP. ■ The value of <code>memberOf</code> attribute sent by the IDP is in the X.500 distinguished format, that is, {cn=groupname,dc=domain}. ■ While adding the IDP configuration to the NetBackup primary server, the values entered for the user group (-g) option matches the SAML attribute name, which is mapped to the <code>memberOf</code> attribute in AD or LDAP. For more information, See “Configure the SAML keystore and add and enable the IDP configuration” on page 505.

Managing role-based access control

This chapter includes the following topics:

- [RBAC features](#)
- [Authorized users](#)
- [Configuring RBAC](#)
- [Default RBAC roles](#)
- [Add a custom RBAC role](#)
- [Role permissions](#)
- [Manage access permission](#)
- [View access definitions](#)

RBAC features

The NetBackup web user interface provides the ability to apply role-based access control in your NetBackup environment. Use RBAC to provide access for the users that do not currently have access to NetBackup. Or, for current NetBackup users with administrator access you can provide limited access and permissions, based on their role in your organization.

For information on access control and auditing information for root users and administrators, refer to the [NetBackup Security and Encryption Guide](#).

Table 35-1 RBAC features

Feature	Description
Roles allow users to perform specific tasks	Add users to one or more default RBAC roles or create custom roles to fit the role of your users. Add a user to the Administrator role to give full NetBackup permissions to that user. See “ Default RBAC roles ” on page 522.
Users can access NetBackup areas and the features that fit their role	RBAC users can perform common tasks for their business role, but are restricted from accessing other NetBackup areas and features. RBAC also controls the assets that users can view or manage.
Auditing of RBAC events	NetBackup audits RBAC events.
DR ready	RBAC settings are protected with the NetBackup catalog.

Authorized users

The following users are authorized to sign in to and use the NetBackup web UI.

Table 35-2 Users that are authorized to use the NetBackup web UI

User	Access	Notes
Root OS administrators Users with the RBAC Administrator role	Full	You can disable automatic access for OS administrators. See “ Disable web UI access for operating system (OS) administrators ” on page 535.
nbasesadmin Appliance user appadmin Flex Appliance user	Default Security Administrator role	This role can grant access to other appliance users. The default admin user for the NetBackup appliance does not have access to the web UI.
Users that have an RBAC role that gives access to the web UI	Varies	

Configuring RBAC

To configure role-based access control for the NetBackup web UI, perform the following steps.

Table 35-3 Steps to configure role-based access control

Step	Action	Description
1	Configure any Active Directory or LDAP domains.	Before you can add domain users, Active Directory or LDAP domains must be authenticated with NetBackup. See the NetBackup Security & Encryption Guide .
2	Determine the permissions that your users need.	Determine the permissions that your users need to perform their daily tasks. You can use the default RBAC roles or use a default role as a template to create a new role. Or, you can create a completely custom role to fit your needs. See “Role permissions” on page 530. See “Default RBAC roles” on page 522. See “Add a custom RBAC role” on page 525.
3	Add users to the appropriate roles.	See “Add a user to a role (non-SAML)” on page 519. See “Add a user to a role (SAML)” on page 521. See “Add a smart card user to a role (non-SAML, without AD/LDAP)” on page 520.
4	Determine the permissions that you want for OS administrators	See “Disable web UI access for operating system (OS) administrators” on page 535. See “Disable command-line (CLI) access for operating system (OS) administrators” on page 534.

Notes for using NetBackup RBAC

Note the following when you configure the permissions for RBAC roles:

- RBAC only controls access to the web UI and not the NetBackup Administration Console.
- When you create roles, be sure to enable the minimal number of permissions so the user can sign in to and use the web UI. Some individual permissions do not have a direct correlation with a screen in the web UI. Users that attempt to sign in but that only have a permission of this kind receive an “Unauthorized” message.
- If a user is added to or removed from a role, the user must sign out and sign in again before the user’s permissions are updated.
- Most permissions are not implicit.

In most cases a **Create** permission does not give a user **View** permission. A **Recovery** permission does not give a user **View** permission or other recovery options like **Overwrite**.

- Not all RBAC-controlled operations can be used from the NetBackup web UI. These types of operations are included in RBAC so a role administrator can create roles for API users as well as for web UI users.
- Some tasks require a user to have permissions in multiple RBAC categories. For example, to establish a trust relationship with a remote primary server, a user must have permissions for both **Remote primary servers** and **Trusted primary servers**.

Add AD or LDAP domains

NetBackup supports Active Directory (AD) or Lightweight Directory Access Protocol (LDAP) domain users. Before you can add domain users to RBAC roles, you must add the AD or the LDAP domain. A domain also must be added before you can configure that domain for smart card authentication.

You can use the `POST /security/domains/vxat` API or the `vssat` command to configure domains.

For more information on the `vssat` command and more of its options, see the [NetBackup Command Reference Guide](#). For troubleshooting information, see the [NetBackup Security & Encryption Guide](#).

View users in RBAC

You can view the users that have been added to RBAC and the roles that they are assigned to. The **Users** list is view-only. To edit the users that are assigned to a role, you must edit the role.

To view the users in RBAC

- 1 On the left, click **Security > RBAC**.
- 2 Click on the **Users** tab.
- 3 The **Roles** column indicates each role to which the user is assigned.

Add a user to a role (non-SAML)

This topic describes how to add a non-SAML user or group to a role.

Non-SAML users use one of the following sign-in methods: **Sign in with username and password** or **Sign in with smart card**.

To add a user to a role (non-SAML)

- 1 On the left, click **Security > RBAC**.
- 2 Click the **Roles** tab.
- 3 Click on the role name, then click on the **Users** tab.
- 4 (Conditional) From the **Sign-in type** list, select from the following:
 - **Default sign-in.** For a user that signs into NetBackup with their username and password.
 - **Smart card user.** For a user that uses a smart card to sign into NetBackup.

Note: The **Sign-in type** list is only available if there is an IDP configuration available for NetBackup.

- 5 Enter the user or the group name that you want to add.

For this type of user	Use this format	Example
Local user or group	<i>username</i>	jane_doe
	<i>groupname</i>	admins
Windows user or group	<i>DOMAINusername</i>	WINDOWS\jane_doe
	<i>DOMAINgroupname</i>	WINDOWS\Admins
UNIX user or group	<i>username@domain</i>	john_doe@unix
	<i>groupname@domain</i>	admins@unix

- 6 Click **Add to list**.
- 7 The user must sign out and sign in again before the user's permissions are updated.

Add a smart card user to a role (non-SAML, without AD/LDAP)

This topic describes how to add a smart card user to a role. In this case the user is a non-SAML user and there is no AD or no LDAP domain association or mapping. User groups are not supported with this type of configuration.

This type of user uses the following sign-in method: **Sign in with smart card**.

To add a smart card user to a role (non-SAML, without AD/LDAP)

- 1 On the left, click **Security > RBAC**.
- 2 Click the **Roles** tab.

- 3 Click on the role name, then click on the **Users** tab.
- 4 (Conditional) From the **Sign-in type** list, select **Smart card user**.

Note: The **Sign-in type** list is available only if there is an IDP configuration available for NetBackup. The smart card user option in the **Sign-in type** list is available when the smart card configuration is done without AD or LDAP domain mapping.

- 5 Enter the username that you want to add.
Provide the exact common name (CN) or the universal principal name (UPN) that is available in the certificate.
- 6 Click **Add to list**.
- 7 The user must sign out and sign in again before the user's permissions are updated.

Add a user to a role (SAML)

This topic describes how to add a SAML user or group to a role.

SAML users use one of the following sign-in methods: **SAML user** or **SAML group**.

To add a user to a role (SAML)

- 1 On the left, click **Security > RBAC**.
- 2 Click the **Roles** tab.
- 3 Click on the role name, then click on the **Users** tab.
- 4 From the **Sign-in type** list, select the sign-in method **SAML user** or **SAML group**.
- 5 Enter the user or the group name that you want to add.

For example, nbuadmin@my.host.com.

If your Identity Provider (IDP) returns group information in the format of (CN=groupname, DC=domainname) or domainname\groupname, you should add the group using the format groupname@domainname. However, it is also possible to configure SAML Groups in Role-Based Access Control (RBAC) without including the domain name. If your IDP returns group names without domain information, you can add those groups as plain text. Please note that using the email format is not mandatory for SAML groups.

- 6 Click **Add to list**.
- 7 The user must sign out and sign in again before the user's permissions are updated.

Remove a user from a role

You can remove a user from a role when you want to remove permissions for that user.

If a user is removed from a role, the user must sign out and sign in again before the user's permissions are updated.

To remove a user from a role

- 1 On the left, click **Security > RBAC**.
- 2 Click the **Roles** tab.
- 3 Click on the role that you want to edit, select the **Users** tab.
- 4 Locate the user you want to remove and click **Actions > Remove > Remove**.

Default RBAC roles

The NetBackup web UI provides the following default RBAC roles with preconfigured permissions and settings.

Table 35-4 Default RBAC roles in the NetBackup web UI

Role name	Description
Administrator	The Administrator role has full permissions for NetBackup and can manage all aspects of NetBackup.
Default AHV Administrator	This role has all the permissions that are necessary to manage Nutanix Acropolis Hypervisor and to back up those assets with protection plans.
Default Apache Cassandra Administrator	This role has all the permissions that are necessary to manage and protect Apache Cassandra assets with protection plans.
Default Cloud Administrator	This role has all the permissions that are necessary to manage cloud assets and to back up those assets with protection plans. Note that a PaaS administrator requires some additional permissions that you can add to a custom role. Cloud administrators also need additional permissions to manage cloud and PaaS assets using intelligent groups. See “Add a custom RBAC role for a PaaS administrator” on page 529.

Table 35-4 Default RBAC roles in the NetBackup web UI (*continued*)

Role name	Description
Default Cloud Object Store Administrator	This role has all the permissions to manage the protection for cloud objects using classic policies.
Default DB2 Administrator	This role provides the ability to view and restore DB2 backups with the <code>nbdb2adutl</code> command. The administrator can also view and manage DB2 jobs.
Default IRE SLP Administrator	Manages IRE (Isolated Recovery Environment) SLP (Storage lifecycle policies) functionalities.
Default Kubernetes Administrator	This role has all the permissions that are necessary to manage Kubernetes and to back up those assets with protection plans. The permissions for this role give a user the ability to view and manage jobs for Kubernetes assets. To view all jobs for this asset type, a user must have the default role for that workload. Or, a similar custom role must have the following option applied when the role is created: Apply selected permissions to all existing and future workload assets .
Default Microsoft Sentinel Administrator	This role has all the permissions necessary to add Microsoft Sentinel credentials in NetBackup and to send NetBackup audit events to Microsoft Sentinel.
Default Microsoft SQL Server Administrator	This role has all the permissions that are necessary to manage SQL Server databases and to back up those assets with protection plans. In addition to this role, the NetBackup user must meet the following requirements: ■ Member of the Windows administrator group. ■ Have the SQL Server “sysadmin” role.
Default Multi-Person Authorization (MPA) Approver	This role has permissions to manage MPA tickets.
Default MySQL Administrator	This role has all the permissions that are necessary to manage MySQL instances and databases and to back up those assets with protection plans.
Default NAS Administrator	This role has all the permissions that are necessary to perform the backup and restore of NAS volumes using a NAS-Data-Protection policy. To view all jobs for the backups and restores of a NAS volume, a user must have this role. Or, the user must have a custom role with same permissions applied when the role was created.
Default NetBackup Command Line (CLI) Administrator	This role has all the permissions that are necessary to manage NetBackup using the NetBackup command line (CLI). With this role a user can run most of the NetBackup commands with a non-root account. Note: A user that has only this role cannot sign into the web UI.
Default Oracle Administrator	This role has all the permissions that are necessary to manage Oracle databases and to back up those assets with protection plans.

Table 35-4 Default RBAC roles in the NetBackup web UI (*continued*)

Role name	Description
Default PostgreSQL Administrator	This role has all the permissions that are necessary to manage PostgreSQL instances and databases and to back up those assets with protection plans.
Default Resiliency Administrator	This role has all the permissions to protect the Veritas Resiliency Platform (VRP) for VMware assets.
Default RHV Administrator	This role has all the permissions that are necessary to manage Red Hat Virtualization computers and to back up those assets with protection plans. This role gives a user the ability to view and manage jobs for RHV assets. To view all jobs for RHV assets, a user must have this role. Or, the user must have a similar custom role with following option applied when the role was created: Apply selected permissions to all existing and future RHV assets.
Default SaaS Administrator	This role has all the permissions to view and manage SaaS assets.
Default Security Administrator	This role has permissions to manage NetBackup security including role-based access control (RBAC), certificates, hosts, identity providers and domains, global security settings, and other permissions. This role can also view settings and assets in most areas of NetBackup: workloads, storage, licensing, and other areas.
Default Storage Administrator	This role has permissions to configure disk-based storage and storage lifecycle policies. SLP settings are managed with the Administrator role.
Default Universal Share Administrator	This role has the permissions to manage policies and storage servers. It can also manage the assets for Windows and Standard client types and for universal shares.
Default Veritas Alta View Administrator	This role has all the permissions that are necessary to manage Cohesity Alta View functionalities.
Default VMware Administrator	This role has all the permissions that are necessary to manage VMware virtual machines and to back up those assets with protection plans. To view all jobs for VMware assets, a user must have this role. Or, the user must have a similar custom role with following option applied when the role was created: Apply selected permissions to all existing and future VMware assets.
NetBackup Read-Only Operator	This role provides the read-only permissions to the IT Analytics Operator, Multi-Person Authorization Approver, and other operators in NetBackup, with no permissions for security.

Note: Cohesity reserves the right to update the RBAC permissions for default roles in future releases. Any revised permissions are automatically applied to users of these roles when NetBackup is upgraded. If you have copies of default roles these roles are not updated automatically. (Or, if you have any custom roles that are based on default roles.) If you want these custom roles to include changes to default roles, you must manually apply the changes or recreate the custom roles.

Add a custom RBAC role

Create a custom RBAC role if you want to manually define the permissions and the access that users have to workload assets, protection plans, or credentials.

Note: Cohesity reserves the right to update the RBAC permissions for default roles in future releases. Any revised permissions are automatically applied to users of these roles when NetBackup is upgraded. Any copies of default roles (or any custom roles that are based on default roles) are not automatically updated.

To add a custom RBAC role

1 On the left, select **Security > RBAC**.

2 Select the **Add** button

3 Select the type of role that you want to create.

You can make a copy of a default role that contains all the preconfigured permissions and settings for that type of role. Or, select **Custom role** to manually configure all the permissions for a role.

4 Provide a **Role name** and a description.

For example, you may want to indicate that the role is for any users that are backup administrators for a particular department or region.

5 Under **Permissions**, select the **Edit** button or the **Assign** button.

The permissions that you select determine the other settings that you can configure for the role.

If you select a default role type, certain permissions are enabled only if they are required for that type of role. (For example, the **Default Storage Administrator** does not require permissions for protection plans. The **Default Microsoft SQL Server Administrator** requires credentials.)

- **Workloads** are enabled when you select **Asset** permissions.
- **Protection plans** are enabled when you select **Protection plans** permissions.

- **Credentials** are enabled when you select **Credentials** permissions.
- 6 Configure the permissions for the role.
See “[Role permissions](#)” on page 530.
See “[Notes for using NetBackup RBAC](#)” on page 518.
- 7 Under **Users**, select the **Assign** button.
- 8 When you are done configuring the role, select the **Save** button.
Note: After a role is created, you must edit permissions for assets, protection plans, or credentials directly from the applicable node in the web UI. For example, to edit permissions for all VMware assets, go to **Workloads > VMware** and then select **VMware settings > Manage permissions**. Or, select a VM and select **Manage permissions**.

Edit or remove a role a custom role

You can edit or remove a custom role when you want to change or remove permissions for users with that role. Default roles cannot be edited or removed. You can only add or remove users from default roles.

Edit a custom role

Note: When you change permissions for a custom role, the changes affect all users that are assigned to that role.

To edit a custom role

- 1 On the left, click **Security > RBAC**.
- 2 On the **Roles** tab, locate and click on the custom role that you want to edit.
- 3 To edit the role description, click **Edit name and description**.
- 4 Edit the permissions for the role. You can edit the following details for a role:

Global permissions for the role

On the **Global permissions** tab, click **Edit**.

Users for the role

Click the **Users** tab.

Access definitions for the role

Click the **Access definitions** tab.

See “[Role permissions](#)” on page 530.

See “[Notes for using NetBackup RBAC](#)” on page 518.

- 5 To add or remove users for the role, click the **Users** tab.
See [“Add a user to a role \(non-SAML\)”](#) on page 519.
See [“Remove a user from a role”](#) on page 522.
- 6 Permissions for assets, protection plans, and credentials must be edited directly in the applicable node in the web UI.

Remove a custom role

Note: When you remove a role, any users that are assigned to that role lose the permissions that the role provided.

To remove a custom role

- 1 On the left, click **Security > RBAC**.
- 2 Click the **Roles** tab.
- 3 Locate the custom role that you want to remove and select the check box for it.
- 4 Click **Remove > Yes**.

Add a custom RBAC role to restore Azure-managed instances

To restore Azure-managed instances, users must have the view permission for these instances. Administrators and similar users can provide other users with a custom role and this permission.

To assign the view permission for Azure-managed instances

- 1 To get the access control ID of the managed instance, enter the following command:

```
GET /asset-service/workloads/cloud/assets?filter=extendedAttributes/managedInstanceName eq 'managedInstanceName'
```

Search for *accessControlId* field in the response. Note down the value of this field.

- 2 To get the role ID, enter the following command:

```
GET /access-control/roles
```

Search for the *id* field in the response. Note down the value of this field.

- 3 Create an access definition, as follows:

```
POST /access-control/managed-objects/{objectId}/access-definitions
```

Request payload

```
{
  "data": {
    "type": "accessDefinition",
    "attributes": {
      "propagation": "OBJECT_AND_CHILDREN"
    },
    "relationships": {
      "role": {
        "data": {
          "id": "<roleId>",
          "type": "accessControlRole"
        }
      },
      "operations": {
        "data": [
          {
            "id": "|OPERATIONS|VIEW|",
            "type": "accessControlOperation"
          }
        ]
      },
      "managedObject": {
        "data": {
          "id": "<objectId>",
          "type": "managedObject"
        }
      }
    }
  }
}
```

Use the following values:

- `objectId`: Use the value of *accessControlId* obtained from step 1.
- `roleId`: Use the value of *id* obtained from step 2.

Note: For an alternate restore, provide the `|OPERATIONS|ASSETS|CLOUD|RESTORE_DESTINATION|` permission in the *operations* list.

Add a custom RBAC role for a PaaS administrator

A PaaS administrator needs additional storage permissions. You can use the **Default Cloud Administrator** role as a template to create a custom role.

To add a custom RBAC role

- 1 On the left, select **Security > RBAC** and click **Add**.
- 2 Select **Default Cloud Administrator**.
- 3 Provide a **Role name** and a description.
For example, you may want to indicate that the role is for any users that are PaaS administrators.
- 4 Under **Permissions**, click **Assign**.
- 5 On the **Global** tab, expand the **Storage** section. Select the following permissions.

Disk pools	View
Storage servers	View
Storage universal shares	View, Create
- 6 On the **Assets** tab, under desired policy type / workload section select the following permissions:
 - Instant access
 - Restore from malware-infected images (Required to restore from malware infected images)
- 7 Click **Assign**.
- 8 Under **Users**, click **Assign**. Then add each user that you want to have access to this custom role.
- 9 When you are done configuring the role, click **Add role**.

Add a custom RBAC role for a Malware administrator

You can use the Default Workload Administrator (of supported workload) role as a template to create a custom role.

To add a custom RBAC role

- 1 On the left, select **Security > RBAC** and click **Add**.
- 2 Select **Default Workload Administrator** or **Custom Role**.
- 3 Provide a **Role name** and a description.
For example, you may want to indicate that the role is for any users that are Malware administrators.
- 4 Under **Permissions**, click **Assign**.
- 5 On the **Global** tab, expand the **NetBackup management** section. Select the following permissions.

Malware	Scan for malware, View scan results
Scan host pools	View, Create, Update, Delete
Scan hosts	View, Create, Update, Delete
Malware tools	View
- 6 Click **Assign**.
- 7 Under **Users**, click **Assign**. Then add each user that you want to have access to this custom role.
- 8 When you are done configuring the role, click **Add role**.

Role permissions

Role permissions define the operations that roles users have permission to perform.

For details on individual RBAC permissions and dependencies, refer to the NetBackup API documentation.

<http://sort.veritas.com>

Table 35-5 Role permissions for NetBackup RBAC

Category	Description
Global	Global permissions apply to all assets or objects. BMR - Configuration and management of BMR. NetBackup Web Management Console Administration - With guidance from Cohesity Support, create diagnostic files to troubleshoot NetBackup and perform JVM garbage collection. These operations are only available from the NetBackup APIs. Refer to the following guides for information on JVM tuning options: NetBackup Installation Guide, NetBackup Upgrade Guide. NetBackup management - Configuration and management of NetBackup. Protection - NetBackup backup policies and storage lifecycle policies. Security - NetBackup security settings. Storage - Manage backup storage settings.
Assets	Manage one or types of assets. For example, VMware assets.
Protection plans	Manage how backups are performed with protection plans.
Credentials	Manage credentials for assets and for other features of NetBackup.

Manage access permission

The **Manage access** permission allows a user to manage who can access a specific part of NetBackup. Users that manage access also need **Access control** permissions. This permission is available for each permission category. However, for some categories the **Manage access** functionality is only available from the NetBackup APIs and not the NetBackup web UI.

For example, a user that has **Manage access** on VMware assets can add or remove the custom roles that have access to VMware assets. This user can also add or remove the specific permissions that a custom role has on VMware assets.

Add the manage access permission to a custom role

If a default role does not have the **Manage access** permission that a user needs, you can create a custom role with that permission. Also give the user the permissions for **User** and **Roles**. These permissions allow the user to view and add users to roles and to add and manage roles.

Assign permissions [Learn about permissions](#)

Global
Assets
Protection plans
Credentials

RHV assets [All](#) | [None](#)

☐ View
☐ Create
☐ Update
☐ Delete

☐ Manage access
☐ Protect
☐ View restore targets
☐ Restore

☐ Allow restore to overwrite
☐ Cancel Jobs
☐ Restart Jobs
☐ View Jobs

VMware assets [All](#) | [None](#)

☒ View
☐ Create
☐ Update
☐ Delete

☒ Manage access
☐ Protect
☐ View restore targets
☐ Restore to cloud

☐ Granular restore
☐ Instant access - Download files
☐ Instant access - Restore files
☐ Instant access

☐ Restore
☐ Allow restore to overwrite
☐ Cancel Jobs
☐ Restart Jobs

☐ View Jobs

Assign permissions

Global
Assets
Protection plans
Credentials

NetBackup management

Protection

Security

Access control

Users

☒ View
☐ Manage access
☒ Assign to role

Roles

☒ Create
☒ Update
☒ Delete
☐ Manage access

Remove access for a custom role

You can remove access to an area of the web UI for a custom role. For each category for which you want to remove the manage access permission, clear the **Manage access** permission. You must edit permissions for assets, protection plans, or credentials directly from the applicable node in the web UI.

For example, to remove manage access permissions for VMware, go to **Workloads > VMware** and then select **VMware settings > Manage permissions**. Or, open the details for a VM and click on the **Permissions** tab.

- 1 On the left, select **Security > RBAC** and click on the **Roles** tab.
- 2 Click on the role.
- 3 Click on the **Access definitions** tab.
- 4 Locate the namespace you want to remove.
- 5 Click **Actions > Remove**.

Disabling access to NetBackup interfaces for OS Administrators

This chapter includes the following topics:

- [Disable command-line \(CLI\) access for operating system \(OS\) administrators](#)
- [Disable web UI access for operating system \(OS\) administrators](#)

Disable command-line (CLI) access for operating system (OS) administrators

By default, an OS administrator (user or group member) has access to the NetBackup CLI and does not need to be a member of an RBAC role.

This option prevents OS administrators from accidentally running NetBackup CLIs. A malicious user with the OS administrator access of the primary server can still bypass this restriction.

After you can disable the option, the OS administrator must log in with `bpnbat -login` to access the CLI.

To disable CLI access for OS administrators

- 1 On the top right, click **Settings > Global security**.
- 2 On the **Security controls** tab, turn off the **CLI access for Operating System Administrator** option.

Disable web UI access for operating system (OS) administrators

By default, an OS administrator (user or group member) has access to the NetBackup web UI and does not need to be a member of an RBAC role.

If you do not want an OS administrator to automatically have this access, you can disable it. An OS administrator must then have the RBAC Administrator role to be able to access the web UI.

To disable web UI access control for the OS administrators

- 1 On the top right, click **Settings > Global security**.
- 2 On the **Security controls** tab, turn off the **Web UI access for Operating System Administrator** option.

Detection and reporting

- [Chapter 37. Detecting anomalies](#)
- [Chapter 38. Malware scanning](#)
- [Chapter 39. Usage reporting and capacity licensing](#)
- [Chapter 40. Reports](#)

Detecting anomalies

This chapter includes the following topics:

- [About backup anomaly detection](#)
- [Configure backup anomaly detection settings](#)
- [View backup anomalies](#)
- [Disable backup anomaly detection and computation of entropy and file attributes for a client](#)
- [About system anomaly detection](#)
- [Configure system anomaly detection settings](#)
- [Configure rules-based anomaly detection](#)
- [Configure risk engine-based anomaly detection](#)
- [View system anomalies](#)

About backup anomaly detection

NetBackup can now detect anomalies in backup metadata. It can detect any unusual job data in the data backup flow. For example, it can detect a file count or a file size that is different than the usual count or size.

Note: By default, the anomaly detection algorithm runs on the NetBackup primary server. If you see any impact on the primary server because of the anomaly detection process, you can configure a media server to detect anomalies.

The following backup job metadata, attributes, or features are verified during backup anomaly detection:

- Backup image size
- Number of backup files
- Data that is transferred in KB
- Deduplication rate
- Backup job completion time

Any unusual deviation in these backup job attributes is considered to be an anomaly and is notified using the NetBackup web UI.

Starting with 10.4, NetBackup can detect anomalies for Oracle workload only for the image size attribute. If an Oracle workload job fails multiple times with status code 5407, NetBackup flags it as an anomaly.

Workflow of backup anomaly detection and notification

The workflow of the backup anomaly detection and notification is as follows:

Table 37-1 Workflow

Step	Description
Step 1	Install or upgrade NetBackup software on the primary server and the media server. See the NetBackup Installation or Upgrade Guide .
Step 2	Enable the primary server to detect backup anomalies. By default, the anomaly detection algorithm runs on the NetBackup primary server. If you see any impact on the primary server because of the anomaly detection process, you can configure a media server to detect anomalies. See the NetBackup Security and Encryption Guide .
Step 3	Configure anomaly detection settings using the NetBackup web UI. See “Configure backup anomaly detection settings” on page 539.
Step 4	View the anomalies using the NetBackup web UI. See “View backup anomalies” on page 542.

How a backup anomaly is detected

Consider the following example:

In an organization, around 1 GB of data is backed up every day for a given client and backup policy with the schedule type FULL. On a particular day, 10 GB of data is backed up. This instance is captured as an image size anomaly and notified. The

anomaly is detected because the current image size (10 GB) is much greater than the usual image size (1 GB).

Significant deviation in the metadata is termed as an anomaly based on its anomaly score.

An anomaly score is calculated based on how far the current data is from the cluster of similar observations of the data in the past. In this example, a cluster is of 1 GB of data backups. You can determine the severity of anomalies based on their scores.

For example:

Anomaly score of Anomaly_A = 7

Anomaly score of Anomaly_B = 2

Conclusion - Anomaly_A is severer than Anomaly_B

NetBackup takes anomaly detection configuration settings (default and advanced if available) into account during anomaly detection.

See the [NetBackup Security and Encryption Guide](#).

Configure backup anomaly detection settings

After you enable anomaly detection, anomaly data gathering, detection service, and events are enabled. Basic and advanced backup anomaly detection settings are available to be configured.

See [“About backup anomaly detection”](#) on page 537.

See [“View backup anomalies”](#) on page 542.

To configure backup anomaly detection settings

- 1 On the left, click **Detection and reporting > Anomaly detection**.
- 2 On the top right, click **Anomaly detection settings > Backup anomaly detection settings**.
- 3 Click **Edit** on the right to configure the following **Anomaly detection > Enable anomaly detection activities** settings:
 - **Enable only for unstructured data** - Enables anomaly detection for the following policy types: Standard, MS-Windows, NAS-Data-Protection, and Universal share.

Note: This is the default configuration for fresh NetBackup 10.4 installation.

- **Enable** - Enables anomaly detection for all policy types except for the ones that are excluded in the **Advanced settings > Disable policy type or specific features for machine learning**.
- **Disable** - Disables anomaly detection in NetBackup for all workload types.
- Click **Save**.

In the case of NetBackup 10.4 upgrade, the value of the **Anomaly detection** option is set based on the previous setting.

- If the option was set to **Enable anomaly data collection, detection service, and events** in the previous version, the option is set to **Enable** after the upgrade.
- If the option was set to a value other than **Enable anomaly data collection, detection service, and events** in the previous version, the option is set to **Disable** after the upgrade.

4 Click **Edit** on the right to configure the **Anomaly detection > Enable automatic scan for imported copy** setting.

- On the **Enable automatic scan for imported copy** pop-up screen, select the **Turn on automatic scan for imported copy** check box.

After enabling the scan for imported copy from the web UI, you must do the following configurations in the `anomaly_config.conf` file:

```
[AUTOMATED_MALWARE_SCAN_SETTINGS]
SCAN_HOST_POOL_NAME=ScanHostPoolName
ENABLE_ALL_CLIENTS=1
TRIGGER_SCAN_FOR_LOW_SEVERITY=1
TRIGGER_SCAN_FOR_MEDIUM_SEVERITY=1
```

- Click **Save**.

5 Select **Edit** to modify the following **Basic Settings**:

- **Anomaly detection sensitivity**
 Use this setting to increase or decrease the sensitivity with which anomalies are detected. If the sensitivity is low, anomalies are detected based on less number of anomalous events.
 If the sensitivity is high, anomalies are detected based on a large number of anomalous events.
- **Data retention settings**
 Use this setting to specify how long you want to retain the anomaly data (in months).
- **Data gathering settings**

Use this setting to specify the time interval (in minutes) after which the anomaly data is gathered for analysis.

- **Anomaly proxy server settings**

Use this setting to specify the NetBackup media server where the anomalies are going to be processed. If not specified, the processing takes place on the primary server.

- Click **Save**.

6 Expand the **Advanced settings** section to configure the following settings:

- Click **Edit** on the right to configure the **Disable anomaly settings for clients** settings.

See [“Disable backup anomaly detection and computation of entropy and file attributes for a client”](#) on page 543.

Click **Save**.

- Click **Edit** on the right to configure the **Disable policy type or specific features for machine learning** settings.

On the pop-up screen, all the policies are listed.

Use the action menus to disable one or all of the following anomaly features for machine learning for the given policy: Backup files count, Data transferred, Deduplication ratio, Image size, and Total time.

- **Disable all** - Use this option to disable all of the anomaly features for machine learning for the given policy.
- **Disable specific features** - Use this option to select specific anomaly features that you want to disable for machine learning.
- Click **Save**.
- Click **Edit** on the right to configure the **Suspicious file extension settings**.
 - Select the **Turn on suspicious file extension detection** to enable NetBackup to detect files with suspicious file extensions.
 A malware such as ransomware attacks the data and encrypts it. After the file encryption, the ransomware renames the files with a specific extension such as `.lockbit`. NetBackup detects such known suspicious file extensions during backups and generates an anomaly.
 - **Files with suspicious extensions (in %)**
 Select the percentage (1 to 50) of files with suspicious extensions from the **Percent** drop-down list, which is acceptable in your environment.
 When the percentage of the files with suspicious extensions exceeds this threshold, an anomaly is generated.
- You can add or remove the suspicious file extensions from the list.

- Click **Save**.

Client offline anomaly type

As part of backup anomaly detection, clients that are offline under suspicious circumstances (with error code 7647) are detected and anomalies are generated.

View backup anomalies

NetBackup can now detect anomalies in backup metadata. It can detect any unusual job data in the data backup flow. For example, it can detect a file count or a file size that is different than the usual count or size.

Consider the following example:

An anomaly of the image size type is displayed as 100MB (Usual 350MB, 450MB). This information implies that the current image size that is reported as anomaly is 100 MB. However, the usual image size range is 350 MB - 450 MB that is derived from the analysis of past data. Because of the significant difference between the current image size and usual image size range, NetBackup notifies it as an anomaly.

See [“About backup anomaly detection”](#) on page 537.

Note: Anomaly count of 0 indicates that there are no anomalies generated or that the anomaly detection services are not running.

To view backup anomalies

- 1 On the left, select **Detection and reporting > Anomaly detection > Backup anomalies**.

The following columns are displayed:

- Job ID - ID of the job for which the anomaly is detected
All child jobs and the associated anomaly details are also shown when you expand the parent job.
- Severity - Severity of the anomalies that are notified for this job
- Asset name - Name of the NetBackup client where the anomaly is detected
- Summary - For the parent job, details like types of anomalies, number of anomalies, and increase or decrease in the number of anomalies are shown. For child jobs, types of anomalies are shown, such as Database corruption.
- Anomaly type - Type of the anomaly such as Image entropy, Job metadata, Suspicious file extension, Client offline

Disable backup anomaly detection and computation of entropy and file attributes for a client

- Backup selection - The backup selection (client or file to be backed up) that is specified in the policy
 - Policy name - The policy name of the associated backup job
 - Policy type - The policy type of the associated backup job
 - Schedule type - The schedule type of the associated backup job
 - Impacted number of jobs - The number jobs for which anomalies are detected
 - Review status - The anomaly status that indicates whether the detected anomaly is reported as a false positive or an actual anomaly, or it can be ignored.
 - Last updated - The date and time when the anomaly status is updated
- 2 Select the job ID to see the job details in the Activity monitor. Expand a parent job to see the details of each child job.
 - 3 You can perform the following actions on the anomaly record:
 - Select **Report as false positive** if the anomaly is a false positive. Similar anomalies are not shown in the future.
The **Review status** of the anomaly record appears as `False positive`.
 - Select **Confirm as anomaly** when you want to take some action on the anomaly condition.
The **Review status** of the anomaly record appears as `Anomaly`.
 - Select **Mark as ignore** when you can ignore the anomaly condition.
The **Review status** of the anomaly record appears as `Ignore`.

Disable backup anomaly detection and computation of entropy and file attributes for a client

You can disable backup anomaly detection and computation of entropy and file attributes for certain NetBackup clients.

To disable backup anomaly detection and computation of entropy and file attributes

- 1 On the left, click **Detection and reporting > Anomaly detection**.
- 2 On the top right, click **Anomaly detection settings > Backup anomaly detection settings**.

- 3 Expand the **Advanced settings** section.
- 4 Select the respective **Edit** option to modify **Disable anomaly settings for clients**.
- 5 On the **Disable anomaly settings for clients** pop-up screen, specify the NetBackup client for which you do not want to generate anomalies and compute entropy.
- 6 In the search results, click the **Add to list** option next to the required client.
- 7 Select **Save**.

The selected clients are added in the excluded clients' list.

Note: After the client is excluded or included again, the computation of entropy and file attributes stops or starts within the next 24 hours with the new backup jobs.

About system anomaly detection

NetBackup can detect system anomalies during critical operations as follows:

- Risk engine-based system anomaly
 See [“Configure risk engine-based anomaly detection”](#) on page 546.
- Monitor database corruption in workloads during job failures - This anomaly monitors database corruption in workloads like MS SQL Server and Oracle during backup job failures.
- Rule-based system anomaly
 See [“Configure rules-based anomaly detection”](#) on page 545.

See [“View system anomalies”](#) on page 550.

Configure system anomaly detection settings

After you enable anomaly detection, anomaly data gathering, detection service, and events are enabled. You can configure specific settings to detect system anomalies in your domain.

See [“About system anomaly detection”](#) on page 544.

To configure system anomaly detection settings

- 1 On the left, click **Detection and reporting > Anomaly detection**.
- 2 On the top right, click **Anomaly detection settings > System anomaly detection configuration**.
- 3 On the **System anomaly detection configuration** screen, configure the following settings:
 - **Risk engine-based anomaly detection**
 See [“Configure risk engine-based anomaly detection”](#) on page 546.
 - **System anomaly detection > Monitor database corruption in workloads during job failures**
 This anomaly monitors database corruption in workloads like Microsoft SQL Server and Oracle during backup job failures.
 Select the checkbox to generate an anomaly alert when NetBackup detects backup job failures because of database corruption in workloads like Microsoft SQL Server and Oracle.
 If database corruption in a workload is detected, status code 5464 is attributed to the parent job that is displayed in the **Activity monitor > Jobs** tab.
 See [“View a job”](#) on page 52.
 Click the status code number to view information about this status code in the Cohesity Knowledge Base.
 See the [NetBackup Status Codes Reference Guide](#).

Note: For detecting database corruption in Microsoft SQL Server, the 'Microsoft SQL Server checksum' option must be set to 'Fail on Error' during MS SQL Server policy configuration.

See [“View backup anomalies”](#) on page 542.

- **Rules-based anomaly detection**
 See [“Configure rules-based anomaly detection”](#) on page 545.

Configure rules-based anomaly detection

Rules engine-based anomaly detection allows you to define certain rules. If the threshold values that are defined in the rule are exceeded, anomalies are generated. For example, an anomaly is generated if a certain number of failed login attempts occur in a specified time period.

For each rule, you can configure the following parameters: execution frequency, query period, and threshold.

To modify the rule parameters, use the `/security/anomaly/rules/{ruleId}` API.

To configure rules-based anomaly detection

- 1 On the left, click **Detection and reporting > Anomaly detection**.
- 2 On the top right, click **Anomaly detection settings > System anomaly detection configuration**.
- 3 On the **System anomaly detection configuration** screen, expand **Rules-based anomaly detection** and select the **Detect anomalies using NetBackup anomaly detection rules** check box.

The following details for each of the predefined rules are displayed:

- Rule name
- Description
- Severity
- Version
- Enabled

For the latest rules file, go to the Cohesity Download Center and download the rules file (.zip) for which you want to generate anomalies.

Select **Upload rules** to select the rules file that you have downloaded. All the latest rules are listed in the **Rules-based anomaly detection** section.

- 4 Select the rules that you want to enable and for which you want to generate anomalies.

Select **Enable**.

NetBackup generates anomalies for the conditions that meet the rule criteria.

Configure risk engine-based anomaly detection

The NetBackup risk engine detects certain system anomalies in a proactive manner and sends appropriate alerts. It helps you take corrective action before you face any security threat in your environment.

You can configure the following options that the risk engine uses to detect anomalies for the given operations:

Detect suspicious image expiration

Use this option to detect when images are expired in an unusual or a suspicious manner.

By default, a system anomaly is generated when the risk engine detects an unusual or a suspicious image expiration attempt and allows the operation to proceed.

However, for additional security, you can configure multi-person authorization for such image expiration attempts, where an MPA approver needs to approve the operation.

Important notes on the Detect suspicious image expiration option

- If the audit retention period is set to less than 3 months, this option accumulates data of 3 months and then becomes active.
- This option supports full backup schedules. Other types of schedule are not considered. The retention level of an image is also not considered for this rule.
- Images are expired by media ID, server name, or by recalculating the retention period.

Select **Edit** and select the **Generate multi-person authorization ticket if images are deleted in a suspicious manner** option.

Note: To successfully review the multi-person authorization tickets, ensure that one or more MPA approvers are available in your environment.

See [“About multi-person authorization”](#) on page 447.

See [“RBAC roles and permissions for multi-person authorization”](#) on page 451.

Detect unusual user sign in

Use this option to detect when a user attempts to sign in to the NetBackup web UI at an unusual time. NetBackup identifies deviations in user sign-in patterns, and flags them.

A notification is generated when an unusual user login is detected.

For additional security, you can configure multi-person authorization for such unusual login attempts, where an MPA approver needs to approve the operation.

See [“Configure multi-person authorization”](#) on page 457.

- If an unusual login attempt is detected on a NetBackup host earlier than 11.0, the request is rejected. Carry out the operation on a NetBackup 11.0 host.
- If an unusual login request is detected in the **NetBackup Administration Console**, the request is rejected. Use the web UI to carry out the operation.

- If none of the users can login and they are placed on hold because of unusual login pattern, the NetBackup administrator can disable the unusual login detection to allow the users to sign in to the NetBackup web UI using the following command:

```
NBU_INSTALL_PATH/netbackup/bin/admincmd/nbsecmd
-disableLoginAnomalyDetection
```

- User logins that are based on the authentication types such as SAML, smart card, and API keys do not support login anomaly detection.

Click **Edit** and use the **Place user's sign in on hold and generate multi-person authorization ticket if a user signs in at an unusual time** option to enable multi-person authentication.

- To successfully review the multi-person authorization tickets, ensure that one or more MPA approvers are available in your environment.
- If multi-person authorization is enabled and an unusual user login is detected, the user's login is placed on hold.
- A ticket is generated and requires approval for the user to proceed. Until the ticket is approved, the user shall not be able to login from the device.
- If the ticket is approved, the user is allowed login and granted a free pass for the next 24 hours. During the free pass period, the user is not subjected to further scrutiny for unusual login attempts.
- If the ticket is rejected, the user cannot log in for the current session but can try again with their credentials.
- The user can choose to cancel their login request.

Detect unusual updates to policies

By default, a system anomaly is generated when the risk engine detects an unusual deletion or update of a policy. An alert is generated and the operation proceeds.

For additional security, you can configure multi-person authorization for such unusual policy update or deletion attempts, where an MPA approver needs to approve the operation.

See [“Configure multi-person authorization”](#) on page 457.

Click **Edit** and use the **Generate multi-person authorization ticket if policy is being modified or deleted in an unusual manner** option to enable multi-person authorization for the **Detect unusual updates to policies** type of anomaly.

- To successfully review the multi-person authorization tickets, ensure that one or more MPA approvers are available in your environment.

- Two alerts are generated for unusual updates to a policy for the next 48 hours. After the second alert, no alert is generated for the next 48 hours even if the policy is modified.
- If multi-person authorization is enabled, a ticket is generated for modification of a policy.
- Approving two consecutive tickets for the same policy does not generate new tickets for the next 48 hours for the same policy.
- If multi-person authorization is enabled for the policy operations on the global level, the **Detect unusual updates to policies** option is disabled.

Note: If multi-person authorization is enabled for the **Detect unusual updates to policies** option, you cannot update or delete policies using the **NetBackup Administration Console** or the command-line interface.

Alternatively, use the `nbcmdrun` command to update or delete policies. For more information on the commands, see the [NetBackup Commands Reference Guide](#).

Secure critical operations

Use this option to protect critical operations such as modifying global security settings and creating API key. When you select this option, you are required to reauthenticate yourself by entering the one-time password that you see in the authenticator application on your smart device before you perform the given critical operations.

Ensure that you have configured multifactor authentication for your user account. If multifactor authentication is not configured, you are not prompted to reauthenticate.

Note: It is strongly recommended that you configure multifactor authentication in your environment to prevent security threats by malicious sources.

See [“Configure multifactor authentication for your user account”](#) on page 468.

Detect possible session hijack

Use this option to detect if there is a possible user session hijack by a malicious source.

The risk engine detects if the same user session token is used by another IP address, and sends a maximum of 10 alerts per day.

Select **Edit** and select the check box to terminate the user session when the risk engine detects that there is a possible session hijack.

View system anomalies

NetBackup can detect system anomalies. This anomaly detection is enabled by default for all policy types.

To view system anomalies

- 1 On the left, select **Detection and reporting > Anomaly detection > System anomalies**.

The following columns are displayed:

- Detected on - The date when the anomaly is detected
- Review status - Anomaly status that indicates whether the detected anomaly is reported as a false positive or an actual anomaly, or it can be ignored.
- Anomaly type - Type of the anomaly
- Severity - Severity of the anomaly
- Description - Additional information about the anomaly
- Anomaly ID - ID of the anomaly record

- 2 Expand a row to see the details of the selected anomaly.

- 3 You can perform the following actions on the anomaly record:

- Select **Report as false positive** if the anomaly is a false positive. Similar anomaly conditions are not reported in the future.
- Select **Confirm as anomaly** when you want to take some action on the anomaly condition.
- Select **Mark as ignore** when you can ignore the anomaly condition.

Malware scanning

This chapter includes the following topics:

- [About malware scanning](#)
- [Configuring a scan host pool](#)
- [Managing a scan host](#)
- [Configure resource limits for malware detection](#)
- [Perform a malware scan](#)
- [Managing scan tasks](#)

About malware scanning

NetBackup finds malware in supported backup images and finds the last good-known image that is malware free. This feature is supported for Standard, MS-Windows, NAS-Data-Protection, Cloud, Cloud-Object-Store, Universal shares, Kubernetes, VMware and Nutanix AHV workload.

NetBackup can scan images that are stored on MSDP, MSDP cloud, AdvancedDisk or OST by using one of the the following **Search by** options:

- **Backup images**

This option is used for scanning policy of client backup images for malware.

- **Assets by policy type**

NetBackup supports MS-Windows, Cloud-Object-Store, NAS-Data-Protection and Standard policy types for malware scan. The following section describes the procedure for scanning NAS-Data-Protection backup images for malware.

NAS-Data-Protection

Each NAS volume or share is read over NFS or SMB, and backed up using a configured number of backup streams. The maximum number of streams per

volume determines the number of backup streams that are created to back up each volume. For example, consider a policy that contains 10 volumes and the maximum number of streams is 4. The backup of the policy creates 4 backup streams for each volume, with a total of 40 child backup streams and 10 parent backup streams.

Note: The number of scans depends on the number of batches that were created to perform the scan. Only the parent stream backup image is visible on the UI.

For more information on multi stream backups, refer to *NetBackup NAS Administrator's Guide*.

■ **Assets by workload type**

This option of malware scanning is used for scanning VMware, Universal shares, Kubernetes, Cloud VM and Nutanix AHV assets for malware.

Note: NetBackup supports VMware assets for malware scan of backup images only with MSDP.

Ensure that the following prerequisites are met:

- The backups were performed with a storage server at NetBackup 10.1 or later.
- Backup images are stored on MSDP storage only with instant access capability, for the supported policy type only.
- The last backup must be successful.
- You must have an RBAC role with permissions to perform malware scans.

Note: According to selection criterion, scan gets initiated to maximum of 100 images.

Benefits of malware scanning

Malware scanning provides the following benefits:

- You can select one or more backup images of the supported policy-types for an on-demand scan. You can use a predefined list of scan hosts.
- If malware is detected during the scanning, a notification is generated in the Web UI.
- In case files are skipped due to not being accessible to scanner or failure from malware scanner, then following respective notifications are generated with information about number and list of skipped files:

- Critical severity: In case malware is found in the backup image and some of the files were skipped during scan.
- Warning severity: In case no malware found in the backup image but some of the files were skipped during scan.

This information can be obtained by clicking on **Actions > Export unscannable files list**.

Note: The malware scan job in **Activity monitor** takes few minutes to reflect the final state of the scan operation running for multiple backup images.

For example, if scan operation runs for 5 backup images in a single request, then the malware scan job in **Activity monitor** would take 5 minutes to reflect the final state which is after completing the last (fifth) backup image scan job.

Note: During recovery if user starts recovery from a malware-affected backup image, a warning message is shown and confirmation is required for proceeding with recovery. Only users with permission to restore from malware-affected images can proceed with recovery.

For more information on best practices for malware scanning, refer to [Smart use of Malware Scanning in NetBackup](#).

Malware scanning before recovery

- User can trigger malware scan of the selected files/folders for recovery as part of recovery flow from Web UI and decide the recovery actions based on malware scan results.
- Catalog entry for the backup image is not updated after recovery time scan as only subset of files are scanned in the backup. Notification would be generated if malware is found as part of recovery time scan.
- During recovery time scan all the images in the start and end date are scanned for malware. Malware scanning of backup image may take long time depending on the number of files selected for recovery. It is recommended to set the Start /End date to include only images which are intended to be used for recovery.
- User can trigger multiple recovery time scan for same backup image.
- Malware scan as part of recovery may take minimum 15-20 minutes for small size backup based on availability of scan host and number of scan jobs in progress. User can track the progress using **Activity monitor > Jobs**. Scan results would be displayed incrementally in the malware detection page. List of backup images in start and end date would be picked up for malware scan incrementally in batches.

- Supported policy types for recovery time scan: Standard, MS-Windows, Cloud-Object-Store, Universal Share, and NAS-Data-Protection.

Note: For successful recovery time malware scan operation, the media server version must be 10.4 or later.


Workflow for malware scanning

This section describes the workflow for malware scanning for the following:

- MSDP backup images using Agentless host as the scan host
- MSDP backup images using NetBackup client as the scan host
- OST and AdvancedDisk

Malware scanning workflow for MSDP backup images using Agentless host as the scan host

The following figure displays the workflow of malware scanning for MSDP backup images:



The following steps depict the workflow for malware scanning for MSDP backup images:

1. After triggering **On Demand Scan**, primary server will validate backup images and create scan jobs for each eligible backup image and identify available scan host for them. Following are few of the criteria's on which the backup images are validated:

- Backup image must be supported for malware detection.
 - Backup image must have a valid Instant Access copy.
 - For an on-demand scan, no existing scan must be running for same backup image. For DNAS the related streams are also considered.
 - Malware detection does not support media server associated with storage.
 - Unable to get information for backup image from catalog.
2. After the backup images are queued for an on-demand scan, the primary server identifies the storage server. An instant access mount is created on the storage server of the configured share type that is specified in scan host pool.

Note: Currently the primary server starts 50 scan threads at a time. After the thread is available it processes the next job in the queue. Until then the queued jobs are in the pending state.

For NetBackup version 10.3 and later, large backups are scanned in batches of 500K files. Each batch is scanned by a separate scan thread.

For recovery time scan, scan in batches feature is not supported.

3. Primary server identifies available and supported MSDP media server and instructs the media server to initiate the malware scan.

If the scan host connectivity type is **Agentless host**, then it instructs the media server to initiate the malware scan.
4. MSDP media server deploys the thin client on the scan host over SSH.
5. Thin client mounts the instant access mount on the scan host.
6. Scan is initiated using the malware tool that is configured in the scan host pool.

Media server fetches the progress of scan from scan host and update the primary server.
7. After the scan is completed, the scan host unmounts the instant access mount from the scan host.
8. Malware scan status is updated to the media server over SSH. Scan logs are copied to the media server log directory.
9. Media server updates the scan status and the infected file list along with skipped file list (if any infected files) to the primary server.
10. Primary server updates the scan results and deletes instant access.
11. Malware scan status notification is generated.

12. Malware scan will timeout in case there is no update on scan. Default timeout period is 48 hours.

Malware detection performs an automated cleanup of eligible scan jobs that are older than 30 days.

Note: The infected scan jobs would be cleaned automatically.

Note: You can download a malware scanner from the Microsoft Azure Marketplace and the AWS Marketplace. Follow the instructions on how to install, configure, and use the malware scanner for AWS and Azure.

Refer to the following for more information:


AWS: [AWS Marketplace](#) and [NetBackup Marketplace Deployment on AWS Cloud](#)

Microsoft Azure: [Microsoft Azure Marketplace](#) and [Microsoft Azure Marketplace](#)

Malware scanning workflow for MSDP backup images using NetBackup client as the scan host

NetBackup version 11.0 and later provides support for **NetBackup client** and **Agentless host** as the scan host to perform the malware scan. The **Agentless host** requires SSH credentials to connect and perform the scan through thin client. The **NetBackup client** uses the client secure communication and performs the scan on the NetBackup client.

The following figure displays the workflow of malware scanning for MSDP backup images:



The following steps depict the workflow for malware scanning for MSDP backup images:

1. After triggering **On Demand Scan**, primary server will validate backup images and create scan jobs for each eligible backup image and identify available scan host for them. Following are few of the criteria's on which the backup images are validated:
 - Backup image must be supported for malware detection.
 - Backup image must have a valid Instant Access copy.
 - For an on-demand scan, no existing scan must be running for same backup image. For DNAS the related streams are also considered.
 - Malware detection does not support media server associated with storage.
 - Catalog must have details of backup image.
2. After the backup images are queued for an on-demand scan, the primary server identifies the storage server. An instant access mount is created on the storage server of the configured share type that is specified in scan host pool.

Note: Currently the primary server starts 50 scan threads at a time. After the thread is available it processes the next job in the queue. Until then the queued jobs are in the pending state.

For NetBackup version 10.3 and later, large backups are scanned in batches of 500K files. Each batch is scanned by a separate scan thread.

For recovery time scan, scan in batches feature is not supported.

To configure the NetBackup client as the scan host, see the *NetBackup Security and Encryption Guide*.

3. Primary server identifies the available and supported MSDP media server and instructs the media server to initiate the malware scan.

If the scan host connectivity type is **NetBackup client**, then the primary server identifies the available NetBackup client scan host from the scan host pool and instructs the NetBackup client scan host to initiate the malware scan.
4. **NetBackup client** as the scan host:
 - NetBackup client mounts the instant access mount on the scan host.
 - Scan is initiated using the malware tool that is configured in the scan host pool.
 - NetBackup client perform the scan operation and updates the progress of scan from scan host to the primary server.

5. After the scan is completed, the scan host unmounts the instant access mount from the scan host.

NetBackup client as the scan host:


- Malware scan status is updated to the primary server. Scan logs are copied to the NetBackup client scan host log directory (`nbmalwarescanner`).
 - NetBackup client scan host updates the scan status and the infected file list along with the skipped file list (if any infected files) to the primary server.
6. Primary server updates the scan results and deletes instant access.
 7. Malware scan status notification is generated.
 8. Malware scan will timeout in case there is no update on scan. Default timeout period is 48 hours.

Malware detection performs an automated cleanup of eligible scan jobs that are older than 30 days.

Malware scanning workflow for OST and AdvancedDisk

For a complete list of supported OpenStorage servers, see the 'OST Storage Servers' section in *NetBackup Hardware and Cloud Storage Compatibility List (HCL)*.

The following figure displays the workflow of malware scanning for OST and AdvancedDisk.



The following prerequisites exist for malware scanning of OST and AdvancedDisk:

- MSDP component for example, SPWS, VPFSD are required for an instant access mount. Hence for OST and AdvancedDisk storage, any one of the media servers must be configured as MSDP storage server so that it can serve the instant access API.

- Primary servers and media servers must be upgraded to NetBackup version 10.4 or later.
- Media servers must be accessible to the OST or AdvancedDisk storage server.
- OST plug-in must be deployed on instant access (host with MSDP components) hosts. For more information on the supported versions of OST plug-ins, refer to the *NetBackup Hardware Compatibility List*.
- Compatible instant access host (RHEL).
- The throttling limit on concurrent instant access from OST and AdvancedDisk STU is same as instant access from MSDP.

The following steps depict the workflow for malware scanning for OST and AdvancedDisk.

1. Using the **On Demand Scan** APIs, the backup image is added to the worklist table on Primary server.

Primary server identifies the available scan host from the specified scan host pool.

2. As part of processing the work list:

(2.1) Create media server for instant access:

- From the backup images, it finds out the storage server.
- From the storage server it finds out the eligible media server.
Media server with instant access capability.
Media server with NetBackup version 10.3 or later.
- Sends the instant access API request to the selected media server.
- If multiple media servers are eligible for an instant access mount request, it selects the media server with minimum number of ongoing instant access requests. This way it can distribute the instant access requests and achieve the load balance.

(2.2) Get IM & TIR

- On the selected media server, in the context of instant access API, it fetches the IM and TIR information from the primary server. It stores the information in the same format that the OS requires for mounting the backup image by VPFSD.
- After instant access mount, for IO file, VPFSD uses OST API to read backup image from storage server.
- Update worklist with images for which instant access was performed with `mountId`, `exportPath`, `storageserver`, and `status`.

3. The primary server identifies the available MSDP media server and instructs the media server to initiate the malware scan.

Note: The media server that is selected for the instant access mount and the server that is selected for communication with the scan host can be the same server or a different server.

4. When it receives the **scan** request, the scan manager from the media server initiates the malware scan on the scan host using thin client (`nbmalwareutil`) through remote communication using SSH.

Note: In NetBackup 10.5 or later, the hash values (SHA-256) of infected files are computed when infected files are found by the NetBackup Malware Scanner. The values can be viewed when exported through **Export infected files list**.

5. Depending on the configuration of scan host, from the scan host it mounts the export using either NFS or SMB from the media server. This media server is where the backup image is mounted using instant access API.
6. Scan is initiated using the malware tool that is configured in the scan host pool.

Note: VPFSD on the media server, uses STS_XXX APIs to open and read the backup images from the OST or AdvancedDisk storage server.

7. After the scan is completed, the scan host unmounts the export path from the media server where backup image is mounted using instant access API.
8. Malware scan status is updated to the media server over SSH. Scan logs are copied to the media server log directory.
9. Media server updates the scan status and the infected file list (if there are any infected files) to the primary server.
10. Primary server updates the scan results and deletes instant access request to the selected media.
11. Malware scan status notification is generated.

Configuring a scan host pool

Refer to the following topics to configure a scan host pool.

See [“Add an existing scan host”](#) on page 563.

See [“Validating the scan host pool configuration”](#) on page 564.

See [“Add a new host in a scan host pool”](#) on page 562.

Prerequisites for scan host pool

Scan host pool is a group of scan hosts. Scan host pool configurations must be performed from NetBackup Web UI before the scan host configuration is completed.

- All the scan host added in the scan host pool must have same malware tool as that of the scan host pool.
- All the scan host added in the pool must have same share type as that of scan host pool.
- *(Applicable only for Agentless scan host)* To add scan host in a scan pool, credentials of scan host and RSA key are required. To get the RSA key of the scan host, refer to the following section:
See [“Managing credentials for malware scanning”](#) on page 565.
- *(Applicable only for NetBackup client as the scan host)* Credentials must be added only for SMB share type. To add the credentials, refer to the following section:
See [“Managing credentials for malware scanning”](#) on page 565.
- Before performing the scan, ensure that the scan hosts are active and available in scan host pool.

Configure a new scan host pool

To configure a new scan host pool

- 1 On left, click **Detection and reporting > Malware detection**.
- 2 On the **Malware detection** page, select **Malware detection settings > Manage malware scanner host pools** on the top-right corner to go to host pool list page.
For configuration details, see the [NetBackup Security and Encryption Guide](#).
- 3 On the **Malware scanner host pools** page, click **Add** to add a new host pool.
- 4 On the **Add scanner host pool** page, enter the details such as **Host pool name**, **Malware scanner**, **Host type**, and **Type of share**.
- 5 Click **Save and add hosts**.

Add a new host in a scan host pool

Use this procedure to add a new scan host in the scan host pool that is configured. See the following topic for prerequisites.

See [“Prerequisites for scan host pool”](#) on page 561.

To add a new host in a scan host pool

- 1 On the left, click **Detection and reporting > Malware detection**.
- 2 Click **Malware detection** page, click **Malware detection settings > Manage malware scanner host pools** on the top-right corner.
- 3 On the **Malware scanner host pools** page, select the desired scan host pool and click **Manage hosts** from the action menu.
- 4 On the **Manage scanner hosts** page, click **Add new**.
- 5 *(Applicable only for Agentless scan host)* On the **Add scanner host to scanner host pool <Host pool name>** page, enter **Host name**.

(Applicable only for NetBackup client as the scan host) On the pop up window, select one or more clients which you want to add.

- 6 *(Applicable only for Agentless scan host)*
 - Click on **Select existing credential** or **Add a new credential**.
See [“Managing credentials for malware scanning”](#) on page 565.
 - Select the media server to validate credentials.
 - Click on **Validate credentials**. On successful validation, click **Save** to save the credentials.
 - Select from the following options:
 - To save the credentials and validate the configuration later, click **Save**.
See [“Validating the scan host pool configuration”](#) on page 564.
 - To save the credentials and immediately validate the configuration, click **Save and validate configuration**.

Note: By default three parallel scans are supported per scan host and this limit is configurable. Having more scan hosts in the scan pool increases the number of parallel scans.

See [“Configure resource limits for malware detection”](#) on page 567.

Managing a scan host

Refer to the following topics to manage a scan host:

- See [“Add an existing scan host”](#) on page 563.
- See [“Validating the scan host pool configuration”](#) on page 564.
- See [“Remove the scan host”](#) on page 564.
- See [“Activate/Deactivate the scan host”](#) on page 565.
- See [“Managing credentials for malware scanning”](#) on page 565.

Add an existing scan host

Use this procedure to add a same scan host in another scan host pool of same share type.

To configure an existing scan host

- 1 On left, click **Detection and reporting > Malware detection**.
- 2 On the **Malware detection** page, click **Malware detection settings > Manage malware scanner host pools** on the top right corner.
- 3 On the **Malware scanner host pools** page, select the desired scan host pool and click **Manage hosts** from the action menu.
- 4 On the **Manage scanner hosts** page, click **Add existing** to select pre-existing host.

Note: List includes all scan hosts from all scan host of agentless or agent based pools.

- 5 On the **Add existing malware scanner host** window, select the desired one or more scan hosts.
- 6 Click **Add**.

Note: (*Applicable only for NetBackup client as the scan host*) After adding the scan host to the scan host pool, user has to manually activate the scan host if it is deactivated.

Validating the scan host pool configuration

(Applicable only for Agentless scan host) Use this procedure to validate the configuration for a new or an existing scan host in the scan host pool that is configured.

To validate the scan host pool configuration

- 1 On the left, click **Detection and reporting > Malware detection**.
- 2 Click **Malware detection** page, click **Malware detection settings** on the top-right corner.
- 3 After adding a new scan host or an existing scan host, on the **Manage scanner hosts** page, select the desired scan host and click **Validate configuration** from the action menu.

To add a new scan host or an existing scan host, see the following topics.

See [“Add a new host in a scan host pool”](#) on page 562.

See [“Add an existing scan host”](#) on page 563.

- 4 On the **Validate configuration** page, enter the details to search and select an image to validate the configuration.

Note: Validating the configuration is only supported for backup image of **Standard** policy type.

- 5 Select the backups to scan and click on **Validate configuration**.

Note: It is recommended to use backup image with small number of files. For large backups, IA creation may delay and test scan might fail.

- 6 After successful validation, click **Finish**.

The **Malware scanner host pools** page is displayed with the list of added scanner hosts.

Remove the scan host

- 1 On left, click **Detection and reporting > Malware detection**.
- 2 On the **Malware detection** page, click **Malware detection settings > Manage malware scanner host pools** on the top right corner.

- 3 On the **Malware scanner host pools** page, select the desired scan host pool and click **Manage hosts** from the Actions menu.
- 4 Select the desired host and click **Delete**, to remove scan host from scan host pool.

Activate/Deactivate the scan host

- 1 On left, click **Detection and reporting > Malware detection**.
- 2 On the **Malware detection** page, click **Malware detection settings > Manage malware scanner host pools** on the top right corner.
- 3 On the **Malware scanner host pools** page, select the desired scan host pool and click **Manage hosts** from the Actions menu.
- 4 Select the desired host and depending on the status of the scan host, click **Activate** or **Deactivate** from the Actions menu.

Managing credentials for malware scanning

To add new credentials

- 1 On the left, click **Detection and reporting > Malware detection**.
- 2 On the **Malware detection** page, select **Malware detection settings > Manage malware scanner host pools** on the top-right corner to go to host pool list page.
- 3 Select the desired scan host. Then click **Actions > Manage hosts**.
- 4 If the scan host connectivity type is **Agentless host**:
 - Select the desired host. Then click **Actions > Manage credentials**.
 - Select **Add a new credential** and click **Next**.
 - Add the details such as the **Credential name**, **Tag**, and **Description**.
 - On the **Host credentials** tab, add the **Host username**, **Host password**, **SSH port**, **RSA key**, and **Share type**.
 - To validate the SSH connection between the MSDP media server and the host, run the following command:
`ssh username@remote_host_name`
 - To verify the RSA key for a remote scan host, run the following command:
`ssh-keyscan scan_host_name 2>/dev/null | grep ssh-rsa`

- To obtain the RSA key for the scan host, use the following command. Use the command from any Linux host with SSH connectivity to the scan host (this host can be the scan host itself):

```
ssh-keyscan scan_host_name 2>/dev/null | grep ssh-rsa | awk
'{print $3}' | base64 -d | sha256sum
```

For example, the output is

33f697637ab3f0911c1d462d4bde8be3eec61a33403e8f6a88daecb415a31eef

- where the RSA key is

33f697637ab3f0911c1d462d4bde8be3eec61a33403e8f6a88daecb415a31eef

Note: Ensure that you remove the - character from the RSA key when you copy.

The following host key algorithms are used to connect to scan host in the given order:

rsa-sha2-512, rsa-sha2-256, ssh-rsa

5 If the scan host connectivity type is **NetBackup client**:

- Select the desired host. Then click **Actions > Manage credentials**.
- Select **Add a new credential** and click **Next**.

To add existing credentials

- 1 On the left, click **Detection and reporting > Malware detection**.
- 2 On the **Malware detection** page, select **Malware detection settings > Manage malware scanner host pools** on the top-right corner to go to host pool list page.
- 3 Select the desired scan host pool. Then click **Actions > Manage hosts**.
- 4 Select the desired host and click **Actions > Manage credentials**.
- 5 Select **Select existing credential**.
- 6 Select the desired credential and click **Select**.

To validate the scan host credentials

(Applicable only if the scan host connectivity type is **Agentless host**)

- 1 On left, click **Detection and reporting > Malware detection**.
- 2 On the **Malware detection** page, select **Malware detection settings > Manage malware scanner host pools** on the top-right corner to go to host pool list page.

- 3 On the **Malware scanner host pools** page, click **Add** to add a new host pool.
The **Add scanner host pool** page appears.
- 4 After you provide the credentials for the scan host on the **Add scanner host pool** page, search for and select the media server.

Note: Only SSH credentials are validated by connecting to scan host from the selected media server. The media server must be a Linux media server with NetBackup version 10.3 or later.

- 5 Click **Validate credential**.
- 6 After the credentials are successfully validated, click **Save**.

Configure resource limits for malware detection

To configure resource limits for malware detection

- 1 On left, click **Detection and reporting > Malware detection**.
- 2 On the top right, click **Malware detection settings > Resource limits**.
For configuration details, see the *NetBackup Security and Encryption Guide*.
- 3 Click **Edit** to edit the resource limit of the resource type.
- 4 Set the global limit which would be considered when resource limit is not set for a resource type.

Or, click **Add** to override the malware global settings.
- 5 Enter the new host name and set the limits.

Note: Resource type scan host: Number of scans per scan host. Default: 3, Minimum: 1, Maximum: 10

Resource type storage server: Number of scans per storage server. Default: 20, Minimum: 1, Maximum: 50

- 6 Click **Save**.

Caution: Setting the Instant Access limit to large value would lead to Storage server resources (memory, CPU, disk) being used for malware scanning purpose. It is advised to set the value based on the existing load on storage server due to backup/duplication operations.

Note: For NetBackup version 10.2 and later, global parallel scans limit configured through **MALWARE_DETECTION_JOBS_PER_SCAN_HOST** configuration option is not applicable. Configure the global parallel scans limit using the Web UI.

Perform a malware scan

To perform a malware scan

- 1 On left, click **Detection and reporting > Malware detection**.
- 2 On the **Malware detection** page, click **Scan for malware**.
- 3 In the **Search by** option, select one of the following:

- **Backup images**
- **Assets by policy type**
- **Assets by workload type**

For more information on the options for scanning, refer to the following on-demand scan:

- See [“Scanning backup images”](#) on page 569.
- See [“Assets by policy type”](#) on page 572.
- See [“Assets by workload type”](#) on page 573.

- 4 Following steps are applicable for scanning **Assets by policy type** and **Assets by workload type**.
 - From the **Client/Asset** table, select a Client/Asset to scan.
 - Click **Next**.

Note: (Applicable only for **Assets by policy type** option) If the selected client in the previous step supports multiple policy types, then user has an option of selecting a single policy type for scanning.

- For the **Start date/time** and **End date/time** verify the date and the time range or update it.
- In the **Scanner host pool**, select the appropriate host pool name.
- (Applicable only for the **NAS-Data-Protection** policy type) In the **Volume** field, select the volume backed up for NAS devices.
Volume-level filtering only fetches the top-level directories of the NAS-Data-Protection volume backup. Volume-level filtering is applicable

only if the top-level directory is a volume. In such a case, you can select individual backup images with the **Backup images** option.

- From the **Current infection status**, select one of the following:
 - **Not scanned**
 - **Not infected**
 - **Infection detected by malware scan**
 - **Infection detected by file hash search**
 - **All**
- Click **Scan for malware**.
 There are more than 100 images in your search. Unable to scan more than 100 images. Adjust the date range and try again.
- After the scan is initiated, the **Scan status** is displayed. Following are the status fields:
 - **Not scanned**
 - **Not infected**
 - **Infected**
 - **Failed**

Note: When we hover on failed status, the tool tip displays the reason for failed scan.

The backup images which failed in validation, are ignored. Malware scanning is supported for the backup images that are stored on storage with instant access capability for the supported policy type only.

- **Pending**
- **In progress**

Scanning backup images

This section describes the procedure for scanning client backup images of a particular policy for malware.

To scan policy of client backup images for malware

- 1** On the left, click **Detection and reporting > Malware detection**.
- 2** On the **Malware detection** page, click **Scan for malware**.

3 In the **Search by** option, select **Backup images**.

4 In the search criteria, review and edit the following:

- **Policy name**

Only supported policy types are listed.

- **Client name**

Displays the clients that have backup images for a supported policy type.

- **Policy type**

Displays all the supported policies which are enabled for malware scanning.

Note: Nutanix-AHV policy would display Nutanix-AHV images, if the backups are taken via Nutanix-AHV policy.

Warning: The **Hypervisor** policy type displays Nutanix AHV and RHV images. NetBackup supports malware scanning only for Nutanix AHV images.

- **Type of backup**

Any incremental backup images that do not have the NetBackup Accelerator feature enabled are not supported for the VMware workload.

- **Copies**

If the selected copy does not support instant access, then the backup image is skipped for the malware scan.

(For *NAS-Data-Protection* policy type) Select the **Copies** as **Copy 2**.

- **Disk pool**

MSDP (PureDisk), OST (DataDomain) and AdvancedDisk storage type disk pools are listed.

- **Disk type**

MSDP (PureDisk), OST (DataDomain) and AdvancedDisk disk types are listed.

- **Infection status**

The malware infected status of the backup images can be searched based on the infection detected by malware scan, file hash search, not infected, not scanned or all.

- For the **Select the timeframe of backups**, verify the date and the time range or update it.

5 Click **Search**.

Select the search criteria and ensure that the selected scan host is active and available.

6 From the **Select the backups to scan** table select one or more images for scan.

7 In the **Select a malware scanner host pool**, **Select** the appropriate host pool name.

Note: Scan host from the selected scan host pool must be able to access the instant access mount created on the storage server which is configured with NFS/SMB share type.

8 Click **Scan for malware**.

9 After the scan is initiated, the **Scan status** is displayed.

The following are the status fields:

- **Not scanned**
- **Not infected**
- **Infected**
- **Failed**

Hover over the status to view the reason for the failed scan.

Note: Any backup images that fail the validation are ignored. Malware scanning is supported for the backup images that are stored on storage with instant access capability and for the supported policy types only.

- **In progress**
- **Pending**

Note: You can cancel the malware scan for one or more in progress and pending jobs.

Assets by policy type

NetBackup supports MS-Windows, Cloud-Object-Store, NAS-Data-Protection and Standard policy types for malware scan. The following section describes the procedure for scanning NAS-Data-Protection backup images for malware.

To scan the supported assets by policy type, perform the following:

- 1** On left, click **Detection and reporting > Malware detection**.
- 2** On the **Malware detection** page, click **Scan for malware**.
- 3** In the **Search by** option, select **Assets by policy type**.
- 4** From the **Client/Asset** table, select a Client/Asset to scan.
- 5** Click **Next**.

If the selected client in the previous step supports multiple policy types, you can select a single policy type for scanning.

- 6** For the **Start date/time** and **End date/time** verify the date and the time range or update it.

The scan is initiated for a maximum of 100 images.
- 7** In the **Scanner host pool**, select the appropriate host pool name.
- 8** In the **Volume** field, **Select volume** backed up for NAS devices.

Note: Volume level filtering only fetches top-level directories of the NAS-Data-Protection volume backup. Volume level filtering is applicable only if the top-level directory is a volume. In such case, user has the option to select individual backup images by using the **Backup images** option in the **Search by** option.

- 9** From the **Current infection status** list select one of the following:
 - **Not scanned**
 - **Not infected**
 - **Infection detected by malware scan**
 - **Infection detected by file hash search**
 - **All**

Note: For NAS-Data-Protection, any backup images that were created on the media server of earlier versions of NetBackup 10.4, select the **All** option for **Current status of malware scan**.

10 Click **Scan for malware**.

Warning: Scan is limited to only 100 images. Adjust the date range and try again.

11 After the scan is initiated, the **Scan status** is displayed. The following are the status fields:

- **Not scanned**
- **Not infected**
- **Infected**
- **Failed**

Note: Hover over the status to view the reason for the failed scan.

Any backup images that fail validation are ignored. Malware scanning is supported for the backup images that are stored on storage with instant access capability and for the supported policy types only.

- **Pending**
- **In progress**

For more information on the malware scan status.

See [“View the malware scan status”](#) on page 574.

Assets by workload type

This section describes the procedure for scanning VMware, Universal shares, Kubernetes, and Cloud VM assets for malware.

This section describes the procedure for scanning VMware, Universal shares, Kubernetes, Nutanix and Cloud VM assets for malware.

To scan the supported assets for malware, perform the following:

1 On left, select the supported workload under **Workloads**.

2 Select the resource which has backups completed.

For example, VMware, Universal shares, Kubernetes, Nutanix and Cloud VM

For example, Nutanix AHV

3 Select **Actions > Scan for malware**.

- 4 On the **Malware scan** page, perform the following:
 - Select the date range for the scan by selecting **Start date/time** and **End date/time**.
 - Select **Scanner host pool**
 - From the **Current infection status** list select one of the following:
 - **Not scanned**
 - **Not infected**
 - **Infection detected by malware scan**
 - **Infection detected by file hash search**
 - **All**

- 5 Click **Scan for malware**.

Note: The malware scanner host can initiate a scan of three images at the same time.

- 6 After the scan starts, you can see the **Scan status** on **Malware detection**, the following fields are visible:
 - **Not scanned**
 - **Not infected**
 - **Infected**
 - **Failed**

Note: Any backup images that fail validation are ignored.

- **In progress**
- **Pending**

Managing scan tasks

View the malware scan status

To view the malware scan status

- ◆ On the left, click **Detection and reporting > Malware detection**.

The following columns are displayed:

- Client - Name of the NetBackup client where the malware is detected.
- Backup time - Time when the backup was performed.
- Scan status - The scan status of the backup image. The different statuses are infected, not infected, failed, in progress, pending, canceled, and cancellation in progress.
- Files infected - Indicates the number of files that were found infected during the scan.
- Scan progress - Indicates the percentage of scan completed.
- Total files - Indicates the count of files and folders as recorded in the catalog for the backup image (list of backup images in case of DNAS). For recovery time scan, the **Total files** column would only indicate the count of files selected for recovery.
- % infected - Provides the percentage of infected files as compared to **Total files**.

Note: Skipped files during recovery are considered as **Not-infected**.

- Elapsed time - Represents the time since scan request was accepted (Date of scan) till the time of completion of scan (End date of Scan). The elapsed time would consist of idle time, time spent in pending state. For resume of failed jobs it would include time spent from failure till the time when the resume operation was triggered.
- Scanned files - Indicates the number of files that are scanned.
- Schedule type - The backup type of the associated backup job
- Date of scan - Date when the scan was performed.
- Policy type - Type of the policy that was selected for scanning.
- Policy name - Name of the policy that was used for scanning.
- Malware scanner - Name of the malware scanner that was used for scanning.
- Scanner host pool - Indicates the host pool used for malware scanning.
- Malware scanner version - Version of the malware scanner that was used for scanning.

Note: To view additional columns that are not displayed, use the **Show or hide columns** pull down menu.

Actions for malware scanned images

Once you scan the backup images for malware detection, a tabular data is available on the **Malware detection** home page.

See [“View the malware scan status”](#) on page 574.

For each backup image, the following quick configurations are available.

To expire all the copies

- 1 On the left, click **Detection and reporting > Malware detection**.
- 2 For the desired scan result, click **Actions > Expire all copies**.
- 3 Confirm to expire all the copies of the selected backup image.

Note: This option is available only for infected scan results.

To view any infected files

- 1 On the left, click **Detection and reporting > Malware detection**.
- 2 For the desired scan result, click **Actions > View infected files**.

Note: This option is available only for infected scan results and scan type 'Recovery'.

- 3 In the **Infected files** table, search or the desired file, if needed.
- 4 If you want to export the list, click **Export list**.

Note: A list of infected files from the selected malware scanning result is exported in `.csv` format. The file name is of the following format:

`backupid_infected_files_timestamp.csv`

To export the infected files list

- 1 On the left, click **Detection and reporting > Malware detection**.
- 2 For the desired malware affected, click **Actions > Export infected files list**.

Note: A `.csv` file contains backup time, names, hashes of the infected files and virus information.

For Microsoft Windows Defender, if real time protection is enabled, then hashes of the infected files are not created as files are not accessible.

To export the unscannable files list

- 1 On the left, click **Detection and reporting > Malware detection**.
- 2 For the desired malware affected, click **Actions > Export unscannable files list**.

Note: A `.csv` file contains the list of files that the malware scanner skips due to issues such as file input or output errors, encrypted (password protected) files, etc.

To cancel a malware scan

- 1 On the left, click **Detection and reporting > Malware detection**.
- 2 For the wanted scan result, click **Actions > Cancel malware scan**.

Note: You can only cancel the malware scan from the "in progress" and the "pending" states.

- 3 Click **Cancel scan** to confirm.

The status changes to **Cancellation in progress**.

Note: The **Cancel malware scan** is not supported for scan results with scan type 'Recovery'.

To rescan an image

- 1 On the left, click **Detection and reporting > Malware detection**.
- 2 For the wanted scan result, click **Actions > Rescan image**.
- 3 Click **Rescan** to confirm.

- 4 For a bulk rescan, when you select one or more images with a different or an empty scanner host pool, you must select a new scanner host pool.
 - Click **Rescan image**.
 - Select a new scan host pool.
The new scan host pool is applicable for all the selected images for this rescan.
 - Click **Rescan** to confirm.
Rescan (and resume) is not supported for scan results with scan type recovery.
- 5 For a rescan of failed or canceled jobs, scanning is triggered from the point of failure (resumed) instead of from a complete scan, under the following conditions:
 - If the value of **Date of scan** is more than 48 hours, then the job is not resumed and the full scan is initiated. This action ensures that the malware signatures that are used for the scan do not differ significantly.
 - Supported for Standard or MS-Windows policy backup images that have a large number of files (> 500 KB). For a DNAS policy, it is supported for more than one stream.
 - Instant Access must have succeeded for the failed job.
 - Resume identifies the first instant access capable copy to scan, which can be different from the copy that was selected for the initial scan request.

After the job is resumed the existing scan result is moved from the state "failed" to "pending" and subsequently to an "in-progress" state. The progress update can continue from the point of failure. For a complete rescan the new scan result is displayed. If the user needs to perform a complete scan, then it can be started using the on-demand scan options.

To delete the scan results

- 1 On the left, click **Detection and reporting > Malware detection**.
- 2 Any scan results that are in a "failed" or "canceled" state can be deleted manually. Click **Actions > Delete scan results**.
- 3 Click **Yes** to confirm the deletion of the selected scan results.
You can select a maximum of 20 scan results to delete.

To view the details of a scan result

- 1 On the left, click **Detection and reporting > Malware detection**.
- 2 Click **Actions > View details** to view details for the backup images with individual batch level.

Note: The **View details** option is available only for the scan results that are in "failed" or "in progress" state.

- 3 On the **View details** page, you can copy information to the Clipboard. Click **Actions > Copy failure details** or **Actions > Copy the scan results**.
- 4 Click **Close**.

Recover from malware-affected images (clients protected by policies)

To restore from malware-affected images, you must have the Administrator role or equivalent RBAC permissions. To recover a VMware asset that is affected by malware, see the following topic.

See [“Recover from malware-affected images \(clients protected by protection plan\)”](#) on page 580.

To recover from malware-affected images (clients protected by policies)

- 1 On the left, click **Recovery**.
- 2 Under **Regular recovery**, click **Start recovery**.
- 3 Select the following properties:

Source client The client that performed the backup.

Destination client The client to which you want to restore the backup.

Policy type The type of policy that is associated with the backup you want to restore.

Restore type The type of restore that you want to perform. The restore types that are available depend on the policy type that you choose.

- 4 Click **Next**.

- 5 Select the **Start date** and **End date**.

Or, click **Use backup history** to view and select specific images. Click **Select** to add the selected images for recovery.

Note: The table displays all the backup image details for selected time frame. You can filter and sort the images based on the malware scan results, infection status, schedule type, policy type or policy name.

- 6 To include any malware-infected images in the recovery, select **Allow the selection of images that are malware-affected**.

Note: The **Allow the selection of images that are malware-affected** option will be disabled if user selects **Scan for malware before recovery** option.

- 7 On the left, expand the **Source client** directory. Select any directories that you want to restore. Or in the right pane, select any files or directories. Click **Next**.
- 8 Select the recovery target.
- 9 To restore any files that are malware-infected, click **Allow recovery of files infected by malware**. Otherwise, NetBackup only restores the files that are scanned and free from malware.
- 10 Select any other recovery options that you want. Then click **Next**.
- 11 Review the recovery settings and then click **Start recovery**.

Recover from malware-affected images (clients protected by protection plan)

To restore from malware-affected recovery points, you must have the Administrator role or equivalent RBAC permissions. To recover a specific recovery point that is affected by malware, see the following topic:

See [“Recover from malware-affected images \(clients protected by policies\)”](#) on page 579.

To recover from malware-affected images for clients protected by protection plan

- 1 On left pane select the supported **Workload**.
- 2 Locate the protected resource and click **Actions > Recover**.
- 3 On the **Recovery points** tab you can see **Malware scan** status of each recovery point, as follows:

- **Not scanned**
 - **Not infected**
 - **Infected**
- 4 Select the recovery point.
 - 5 Select **Allow the selection of recovery points that are malware-affected**. This option only displays if there are recovery points that contain malware-affected images.

Note: To restore from malware-affected recovery points, you must have the Administrator role or equivalent RBAC permissions.

- 6 Click **Recover** and select the type of recovery. Then follow the prompts.
For more details on recovering a VM, see the *NetBackup for VMware Administrator's Guide*.

Clean file recovery for virtual workload (VMware)

To restore from malware-affected images, you must have the Administrator role or equivalent RBAC permissions.

To recover a VMware asset that is affected by malware, refer to the following procedure:

VMware single file restore from recovery point (with agent and agentless)

- 1 On left pane select **Workload > VMware**.
- 2 Search and click on the virtual machine to be recovered.
- 3 On the **Recovery points** tab, select the date for recovery point.

- 4
- Select **Allow the selection of recovery points that are malware-affected**. This option only displays if there are recovery points that contain malware-affected images.

Note: To restore from malware-affected recovery points, you must have the Administrator role or equivalent RBAC permissions.

- 5
- Click **Recover** and select the type of recovery as **Restore files and folders**. Then follow the prompts.

Note: NetBackup now provides support for VMware single file restore clean recovery by selecting the **Allow recovery of files infected by malware** option in the **Recovery options**. This option overrides the default behavior.

For more details on recovering a VM, see the *NetBackup for VMware Administrator's Guide*.

To recover a specific recovery point that is affected by malware, refer to the following procedure:

Single file restore using recovery flow (with agent)

- 1
- On the left, click **Recovery**.
- 2
- Under **Regular recovery**, click **Start recovery**.
- 3
- Select the following properties:

Policy type	The type of policy that is associated with the backup you want to restore. Select the policy type as VMware.
Source client	The client that performed the backup. Under the Virtual machines search tab, select the virtual machine and click Apply .
Destination client	The client to which you want to restore the backup.
Restore type	The type of restore that you want to perform. The restore types that are available depend on the policy type that you choose. Note: Clean recovery is supported only for normal backups.

- 4
- Click **Next**.

5 Edit the **Date range**.

Or, click **Use backup history** to view and select specific images. Click **Apply** to add the selected images for recovery.

Note: The table displays all the backup image details for selected time frame. You can filter and sort the images based on the malware scan results, infection status, schedule type, policy type or policy name.

- 6** To include any malware-infected images in the recovery, select **Allow the selection of images that are malware-affected**.
- 7** On the left, expand the **Source client** directory. Select any directories that you want to restore. Or in the right pane, select any files or directories. Click **Next**.
- 8** Select the recovery target.
- 9** To restore any files that are malware-infected, click **Allow recovery of files infected by malware**. Otherwise, NetBackup only restores the files that are scanned and free from malware.
- 10** Select any other recovery options that you want. Then click **Next**.
- 11** Review the recovery settings and then click **Start recovery**.

Usage reporting and capacity licensing

This chapter includes the following topics:

- [Track protected data size on your primary servers](#)
- [Add a local primary server](#)
- [View license types in usage reporting](#)
- [Download usage reports](#)
- [Scheduling reports for capacity licensing](#)
- [Other configuration for incremental reporting](#)
- [Troubleshooting failures for usage reporting and incremental reporting](#)

Track protected data size on your primary servers

The Usage reporting application displays the primary servers that are configured for capacity licensing and their respective consumption details. This reporting provides the following benefits:

- Ability to plan for capacity licensing.
- On a weekly basis, NetBackup gathers and reports usage and trend information. The `nbdeployutil` utility has scheduled runs to gather data for the report (enabled by default).
- A link to the [Cohesity NetInsights Console](#). The Usage Insights tool within the NetInsights Console tool NetBackup customers to proactively manage their license use through near real-time visibility of consumption patterns.

- Reporting is done for all policy types that are used for data protection.

Requirements

NetBackup automatically collects data for the usage reporting, provided the following requirements are met:

- The primary servers (or primary servers) are at NetBackup 8.1.2 or later.
- You use capacity licensing.
- You use automatic, scheduled reports. If you manually generate capacity license reports, the data does not display in the usage report in the NetBackup web UI.
- The following file exists:

UNIX: `/usr/openv/var/global/incremental/Capacity_Trend.out`

Windows: `install_path\var\global\incremental\Capacity_Trend.out`

The **Usage** tab displays an error if the backup data is not available. Or, if the usage report is not generated (file does not exist).

- If you want one of your primary servers to gather usage reporting data for other remote primary servers, additional configuration is required. You must create a trust relationship between the primary servers. You must also add the local primary server (where you plan to run `nbdeployutil`) to the **Servers** list on each remote primary server.

See [“Add a local primary server”](#) on page 585.

See [“Add a trusted primary server”](#) on page 482.

Additional information

- Details are available on capacity licensing, scheduling, and options for capacity licensing reports.
See [“Scheduling reports for capacity licensing”](#) on page 587.
- *Cohesity Usage Insights for NetBackup Getting Started Guide*. Details on how to use [Usage Insights](#) to manage your NetBackup deployment and licensing. This tool provides accurate, near real-time reporting for the total amount of data that is backed up.

Add a local primary server

If you want to add usage reporting information for a primary server but that server does not have an internet connection, you need to add the name of the local primary server to the servers list of the remote primary server. The local primary server is where you plan to run the usage reporting tool.

To add a local primary server

- 1 On the left, click **Hosts > Host Properties**.
- 2 Select the host and click **Connect**.
- 3 Click **Edit primary server**.
- 4 Click **Servers**.
- 5 On the **Additional Servers** tab, click **Add**.
- 6 Enter the name of the primary server where you plan to run `nbdeployutil`.
- 7 Click **Add**.

View license types in usage reporting

You can view the license types for which you want to generate usage reports using the `netbackup_deployment_insights` utility.

To select license types to display in usage reporting

- 1 On the left, click **Detection and reporting > Usage**.
- 2 On the top right, click **Usage reporting settings**.

The usage reporting type and license model are displayed for the primary server.

Download usage reports

You can download the reports that are automatically generated using the `netbackup_deployment_insights` utility.

To download usage reports

- 1** On the left, click **Detection and reporting > Usage**.
- 2** Click **Download reports**.

The **Download reports** pop-up displays the Excel and JSON files. For each file, the name, the date when the report is generated, and the frequency at which the `netbackup_deployment_insights` utility runs and generates the reports is displayed.

Note: The pop-up may display an Excel file of a newer date and a JSON file of an earlier date. After the telemetry agent cycle is completed, the latest JSON file is displayed.


- 3** Select the report that you want to download and click **Download**.

In an upgrade scenario, the download reports feature is only available after the next successful incremental run of the `netbackup_deployment_insights` utility. If you try to download any NetBackup 10.4.0.1 or earlier reports, the download may fail.

Scheduling reports for capacity licensing

By default, NetBackup triggers `nbdeployutil` to run on a specified schedule to incrementally gather data and to generate licensing reports. For the first run, the duration of the report uses the frequency that is specified in the configuration file.

For capacity licensing, the report duration is always for the last 90 days based on the availability of the gathered data. Any data older than 90 days is not considered in the report. Each time `nbdeployutil` runs, it gathers information for the time between the latest run of `nbdeployutil` and the previous successful run.

Figure 39-1 Generating incremental capacity licensing reports

Licensing report location

The current capacity licensing report resides in the following directory:

On Windows: `install_path\NetBackup\var\global\incremental`

On UNIX: `/usr/openv/var/global/incremental`

It contains the following files:

- The generated report for the latest `nbdeployutil` result.
- Folders containing incrementally gathered data.
- The archive folder that contains the older generated reports.
- `nbdeployutil` log files.

The older reports are placed in the archive folder. Cohesity recommends that you retain at least 90 days of reporting data. Data can be kept longer than 90 days, depending on the requirements of your environment. Older reports can help to show how the capacity usage has changed over time. Delete the reports or the folder when they are no longer required.

From the NetBackup web UI, you can download the reports that are automatically generated using the `nbdeployutil` utility. On the web UI, click **Detection and reporting > Usage > Download reports**. For more information, see the *NetBackup Web UI Administrator's Guide*.

The **Download reports** feature requires that you have sufficient permissions on the gather directory. For the custom path specified in the `PARENTDIR` configuration setting, ensure that the required read permissions are provided for the NetBackup Web service user. If you delete the default incremental folder and then manually create it, ensure that the required read permissions are provided for the NetBackup Web service user.

Use Case I: Using the default values for the licensing report

The `nbdeployutilconfig.txt` file is not required when you use the default parameters. `nbdeployutil` uses the following default values for capacity licensing:

- `FREQUENCY_IN_DAYS=7`
- `MASTER_SERVERS=local_server`
- `PARENTDIR=folder_name`
For Windows: `install_path\NetBackup\var\global\incremental`
For UNIX: `/usr/opensv/var/global/incremental`
- `PURGE_INTERVAL=120` (number of days)
- `MACHINE_TYPE_REQUERY_INTERVAL = 90` (number of days)

Use Case II: Using custom values for the licensing report

If the file `nbdeployutilconfig.txt` is not present, create a file using the following format:

```
[NBDEPLOYUTIL_INCREMENTAL]
MASTER_SERVERS=<server_names>
FREQUENCY_IN_DAYS=7
PARENTDIR=<folder_name_with_path>
PURGE_INTERVAL=120
MACHINE_TYPE_REQUERY_INTERVAL=90
```

To use custom values for the licensing report

- 1 Copy the `nbdeployutilconfig.txt` file to the following location:
For Windows: `install_path\NetBackup\var\global`
For UNIX: `/usr/opensv/var/global`
- 2 Open the `nbdeployutilconfig.txt` file.

- 3 Edit the `FREQUENCY_IN_DAYS` value to reflect how often you want the report to be created.

Default 7
(recommended)

Minimum 1

Parameter deleted `nbdeployutil` uses the default value.

Note: If you enter a value that is less than 1, `nbdeployutil` automatically uses 7, which is the default value.

- 4 Edit the `MASTER_SERVERS` value to include a comma-separated list of the primary servers you want to include in the report.

Note: Cohesity Usage Insights requires that primary servers be at NetBackup 8.1.2 or later.

No value `nbdeployutil` uses the default value.

Parameter deleted `nbdeployutil` uses the default value.

For example:

- `MASTER_SERVERS=newserver, oldserver`
- `MASTER_SERVERS=newserver, oldserver.domain.com`
- `MASTER_SERVERS=myserver1.somedomain.com, newserver.domain.com`

- 5 Edit the `PARENTDIR` value to include the full path for location where the data is gathered and reported.

No value `nbdeployutil` uses the default value.

Parameter deleted `nbdeployutil` uses the default value.

- 6 Edit the `PURGE_INTERVAL` to indicate the interval (in days) for how often you want to delete the report data. Data that is older than 120 days is automatically purged.

Default	120
Minimum	90
No value	<code>nbdeployutil</code> uses the default value.
Parameter deleted	<code>nbdeployutil</code> uses the default value.

- 7 Edit the `MACHINE_TYPE_REQUERY_INTERVAL` to indicate how often to scan physical clients for updates to the machine type.

Default	90
Minimum	1
No value	<code>nbdeployutil</code> uses the default value.
Parameter deleted	<code>nbdeployutil</code> uses the default value.

Other configuration for incremental reporting

To change the directory of the gathered data and capacity licensing report

- 1 If you have older gathered data and licensing reports, copy the complete directory to the new location.
- 2 Edit `nbdeployutilconfig.txt` and change the location of the gathered data and licensing report in the `PARENTDIR=folder_name` field.

To use the data that was gathered previously to generate a capacity licensing report

- 1 Locate the folder that was generated for the gathered data after the previous run of `nbdeployutil` and copy it to the following location:

On Windows: `install_path\NetBackup\var\global\incremental`

On UNIX: `/usr/openv/var/global/incremental`

- 2 Create the `gather_end.json` file inside the copied folder and add the following text:

```
{"success":0}
```

The next incremental run considers the data inside the copied folder to generate a capacity licensing report.

Note: Delete any other gather folders inside the copied folder to avoid gaps for the period in which data is gathered. The missing data is automatically generated during the next incremental run.

To create a custom interval report using existing gathered data for capacity licensing

- ◆ To create a report for a time interval that is different than the default interval of 90 days, run the following command:

On Windows:

```
nbdeployutil.exe --capacity --incremental --report --inc-settings  
"install_dir\netbackup\var\global\nbdeployutilconfig.txt"  
--hoursago <custom-time-interval>
```

On UNIX:

```
nbdeployutil.exe --capacity --incremental --report --inc-settings  
"/usr/openv/var/global/nbdeployutilconfig.txt"  
--hoursago <custom-time-interval>
```

The number of hours specified in `--hoursago` must be fewer than the `purge-interval` that is specified in the `nbdeployutilconfig.txt` file.

You can also use `--start` or `--end` options in the `nbdeployutilconfig.txt` file.

```
--start="mm/dd/yyyy HH:MM:SS"
```

```
--end="mm/dd/yyyy HH:MM:SS"
```

If the latest gather operation fails to retrieve front-end data size (FEDS) data, the custom report fails because the required backup information is not available. Let the next scheduled incremental gather run successfully and then try to generate the custom report.

Note: `nbdeployutil` uses existing gathered data to generate the custom interval report. You are not required to use the `--gather` option.

Troubleshooting failures for usage reporting and incremental reporting

- For incremental runs of `nbdeployutil`, notifications are sent to the NetBackup web UI. The notification details include, status of the run, duration, start time, and end time.

- `nbdeployutil` fails to gather data and generate the report for your environment. Refer to the logs to understand when the task failed and the reason for the failure.
- `nbdeployutil` fails with a `bpimagelist` error with status 37 after you run the utility manually. Ensure that you added the primary servers to the additional servers list.
- The following error displays because of internal web service communication failures:
Internal Web API error occurred for primary server *SERVER_NAME*. Run `nbdeployutil` again with the gather option on primary server *SERVER_NAME*.
- For VMware or NDMP, when the backup agent fails to post licensing information to the database, a status code 5930 or 26 displays in the Activity Monitor: For more information, see the [NetBackup Status Codes Reference Guide](#).
- `nbdeployutil` may fail with errors related to loading the Perl modules. In such a scenario, it is recommended to refer the Perl documentation related to the reported error.

You can use `netbackup_deployment_insights` with the same troubleshooting points.

Reports

This chapter includes the following topics:

- [About the reports utility](#)
- [Run a report](#)
- [Copy a report text to another document](#)

About the reports utility

Use the Reports utility to generate reports to verify, manage, and troubleshoot NetBackup operations. NetBackup reports display information according to job status, client, backups, and media contents.

NetBackup offers many different reports to view information about job activity and media. In NetBackup 11.0, only the **All log entries** report is available in the NetBackup web UI. Other reports are available in the NetBackup Administration Console. See the [NetBackup Administrator's Guide, Volume I](#) for details.

Table 40-1 Reports

All log entries	The All Log Entries report generates a list of all log entries for the specified time period.

Run a report

The following procedure describes how to run a NetBackup report from the **Reports** utility.

To run a report

- 1 In the NetBackup web UI, on the left select **Detection and reporting > Reporting**.
- 2 Select the criteria for what to include or exclude in the report. For example:
 - Select the time period that the report should span.
 - Select the media servers and clients on which to run the report.
 - Enter a **Job ID** to run reports for that Job ID.
- 3 Select **Run report**.

Copy a report text to another document

The following procedure describes how to copy the text from a NetBackup report and paste it into a spreadsheet or other document.

To copy report text to another document

- 1 In the NetBackup web UI, on the left select **Detection and reporting > Reporting**.
- 2 Select the criteria for what to include or exclude in the report, and select **Run report**.
- 3 Select the rows of the report you want to copy by holding down the **Shift** or **Ctrl** key.
- 4 Select **Copy to clipboard**.
- 5 Paste the selected rows into a spreadsheet or other document.

NetBackup workloads and NetBackup Flex Scale

- [Chapter 41. NetBackup SaaS Protection](#)
- [Chapter 42. NetBackup Flex Scale](#)
- [Chapter 43. NetBackup workloads](#)

NetBackup SaaS Protection

This chapter includes the following topics:

- [Overview of NetBackup for SaaS](#)
- [Adding NetBackup SaaS Protection Hubs](#)
- [Configuring the autodiscovery frequency](#)
- [Viewing asset details](#)
- [Configuring permissions](#)
- [Troubleshooting SaaS workload issues](#)

Overview of NetBackup for SaaS

The NetBackup web UI provides the capability to view the assets of NetBackup SaaS Protection. The assets configured to protect SaaS applications data are automatically discovered in the NetBackup web UI.

The NetBackup SaaS Protection assets comprise of the assets such as, Hubs, StorSites, Stors, and Services.

The following assets details are displayed:

- Storage size
- Storage tier details
- Number of items in the storage
- WORM details
- Write, delete, stub policies details

- Schedule for the next backup
- Status of the last backup

The NetBackup web UI lets you perform the following operations:

- Add a NetBackup SaaS Protection Hub.
- View assets in the Hub.
- Launch the NetBackup SaaS Protection web UI.
- Delete the added Hub.

Note: If a SaaS asset is deleted from NetBackup SaaS Protection web UI, the deleted asset is not removed from the NetBackup database immediately. The deleted asset remains in the NetBackup database for 30 days.

The following table describes the features of NetBackup for SaaS:

Table 41-1 Features of NetBackup for SaaS

Features	Description
Integration with NetBackup role-based access control (RBAC)	The NetBackup web UI provides RBAC roles, which enable its users to view assets in SaaS workload. The user does not need to be a NetBackup administrator to add a NetBackup SaaS Protection Hub or view assets in the Hub.
NetBackup SaaS Protection-specific credentials	NetBackup SaaS Protection service accounts are used to authenticate the Hubs.
Autodiscovery of assets	NetBackup automatically discovers StorSites, Stors, and Services in the Hubs. You can also perform manual discovery. After the assets are discovered, you can view assets details.
Cross Launch	You can cross launch the NetBackup SaaS Protection web UI. If SSO is configured the user is redirected to the NetBackup SaaS Protection UI without entering the credentials for each login.

About NetBackup SaaS Protection

NetBackup SaaS Protection is a cloud-based data protection solution that is deployed on Microsoft Azure. It is used to protect the data of the on-premises application and SaaS applications.

NetBackup SaaS Protection protects data of the following SaaS applications:

- Box
- Exchange
- Google Drive
- SharePoint sites
- OneDrive sites
- Teams sites and chats
- Slack

NetBackup SaaS Protection supports bulk or granular data restore at the required locations. It also supports restore of the last updated data or any specific point-in-time data.

An account is configured for a customer, which is referred to as a tenant. The assets are configured for the tenant to protect the required data.

For more information, refer to the NetBackup SaaS Protection administrator's guide.

Adding NetBackup SaaS Protection Hubs

You can add NetBackup SaaS Protection Hubs and autodiscover all the assets inside the Hub.

To add NetBackup SaaS Protection Hubs

- 1 On the left, click **Workloads > SaaS**.
- 2 On the **Hubs** tab, click **Add**.
- 3 On the **Add a NetBackup SaaS Protection Hub** page, enter the name of the Hub.
 - To use the existing credential, click **Select existing credentials**.
On the next page, select the required credentials, and click **Select**.
 - To create a new credential, click **Add a new credential**.
On the **Add credential** page, enter the following:
 - **Credentials name**: Enter a name for the credential.
 - **Tag**: Enter a tag to associate with the credential.
 - **Description**: Enter a description of the credential.
 - **Username**: Enter the username, which is configured as a service account in NetBackup SaaS Protection.

- **Password:** Enter the password.
- 4 Click **Add**.
After the credentials are successfully validated, the Hub is added and autodiscovery runs to discover available assets in the Hub.
See [“Configure NetBackup for single sign-on \(SSO\)”](#) on page 501.

Configuring the autodiscovery frequency

Autodiscovery keeps a count of the assets in Hubs. NetBackup web UI refreshes Hubs at intervals to get the updates from NetBackup SaaS Protection for any addition or removal of assets. By default, the interval for refresh is 8 hours.

To configure the autodiscovery frequency

- 1 On the left, click **Workloads > SaaS**.
- 2 On the top-right, click **SaaS settings > Autodiscovery**.
- 3 Click **Edit**.
- 4 Enter the number of hours after which NetBackup should run autodiscovery and click **Save**.

Viewing asset details

The NetBackup SaaS Protection assets are displayed in two tabs, **Services** tab and **Hubs** tab.

To view asset details

- 1 On the left, click **Workloads > SaaS**.
The **Services** tab is displayed. It displays the services configured for the Hub.
You can perform the following actions on the tab:
 - View the list of services configured for the Hub.
 - Search the required service in the list of services.
 - Filter the list of services based on the status of the services.
 - Sort columns.
 - View the following service details:
 - Application type for which the service is configured.
 - Date and time of the last backup and the next scheduled backup.

- Criteria set for write, stub, and delete policy.
- WORM details.

2 Click the **Hubs** tab to view details on Hubs, StorSites, and Stors.

You can navigate to the required asset using the left panel. You can perform the following actions on the **Hubs** tab:

- View a list of the Hubs.
- Search for a Hub in the list.
- Add new Hubs.
- Validate the credentials.
- Sort columns.
- Click **Actions** to perform the following:
 - Edit credentials.
 - Delete the Hub.
 - Manually discover assets in the Hub.
- View the following asset details:
 - Associated Stors, last backup details, and so on for the Services.
 - Version, ID, and state of the Hub.
 - State, tier details, and so on for the StorSite.
 - State, policy details, and so on for the Stors.
 - Launch the NetBackup SaaS Protection web UI. You can cross launch the NetBackup SaaS Protection web UI from Services, Stors, and Hubs page.

For more information, refer to the NetBackup SaaS Protection administrator's guide.

Configuring permissions

Using the NetBackup web UI, you can assign different access privileges to the user roles on the assets. For example, view, update, delete, and manage access.

See [“Manage access permission”](#) on page 531.

Note: The user with access permission on the SaaS workload in NetBackup, and no or limited permissions in NetBackup SaaS Protection can still view the NetBackup SaaS Protection assets on the NetBackup web UI.

Troubleshooting SaaS workload issues

Check the following locations for logs of the SaaS workload:

- PiSaaS
 - Windows: <install path>\Veritas\NetBackup\logs\ncfnbcs
 - UNIX: <install path>/openv/netbackup/logs/ncfnbcs
- bpVMUtil
 - Windows: <install path>\Veritas\NetBackup\logs\bpVMutil
 - UNIX: <install path>/openv/netbackup/logs/bpVMutil
- APIs/nbWebServices
 - Windows: <install path>\Veritas\NetBackup\logs\nbwebsevice
 - UNIX: <install path>/openv/logs/nbwebsevice

Use the following information to troubleshoot issues.

Table 41-2 Troubleshooting issues in SaaS Workload

Problems	Recommended actions
Failed to add a Hub due to incorrect Hub name or invalid user credentials.	Enter appropriate Hub name and valid credentials.
Failed to add a Hub due to issue in credential validation.	Check if the credentials are not expired. Also check if the credentials are valid.
Failed to add a Hub due to limited permissions.	Assign appropriate permissions to the user on the SaaS workload. See “Role permissions” on page 530.
Failed to delete a Hub due to limited permissions.	Assign appropriate permissions to the user on the SaaS workload. See “Role permissions” on page 530.

Table 41-2 Troubleshooting issues in SaaS Workload (*continued*)

Problems	Recommended actions
Failed to perform discovery on the Hub due to limited permissions.	Assign appropriate permissions to the user on the SaaS workload. See "Role permissions" on page 530.
The services are not deleted from NetBackup after deleted the associated Connector from NetBackup SaaS Protection.	The services get removed from NetBackup after 30 days from Connector deletion.
Failed to launch the NetBackup SaaS Protection web UI using the Launch NSP option. Credentials are required while launching the NetBackup SaaS Protection web UI.	Check if SSO is configured correctly. If SSO is configured correctly, check if the user has appropriate permissions to access the NetBackup SaaS Protection web UI. See "Configure NetBackup for single sign-on (SSO)" on page 501.
Connecting to the proxy host X.X.X.X on port 3128 with type SOCKS5	Configure proxy settings on the primary server using the bpsetconfig utility.

NetBackup Flex Scale

This chapter includes the following topics:

- [Managing NetBackup Flex Scale](#)

Managing NetBackup Flex Scale

The NetBackup Flex Scale appliance administrator can access Cluster Management in the NetBackup web UI. The Appliance administrator must be assigned the RBAC **Administrator** role for the NetBackup web UI.

For full instructions on managing NetBackup Flex Scale, see the following resources.

NetBackup Flex Scale Installation and Configuration Guide

NetBackup Flex Scale administrator's guide

Table 42-1 Accessing NetBackup Flex Scale and NetBackup


Interface and URL	Access to NetBackup Flex Scale or NetBackup
NetBackup web UI https://primaryserver/webui/login	To open NetBackup Flex Scale, click the Appliance management node. This action opens the NetBackup Flex Scale infrastructure management console in a new browser tab. See “Access NetBackup Flex Scale from the NetBackup web UI” on page 608.
NetBackup Flex Scale web UI https://ManagementServerIPorFQDN/webui	To access the NetBackup Flex Scale features, expand Cluster Management . See “Manage NetBackup and the NetBackup Flex Scale cluster management from the NetBackup Flex Scale web UI” on page 607.

Table 42-1 Accessing NetBackup Flex Scale and NetBackup (*continued*)

Interface and URL	Access to NetBackup Flex Scale or NetBackup
NetBackup Flex Scale infrastructure management console IPv4: <code>https://ManagementServerIPorFQDN:14161/</code> IPv6: <code>https://ManagementServerIP:14161/</code>	To open NetBackup, click the NetBackup node. This action launches the NetBackup Flex Scale web UI in the same browser tab. To access the NetBackup Flex Scale infrastructure management console again, click Cluster Management. See “Access NetBackup from the Flex Scale infrastructure management console” on page 606.

Access NetBackup from the Flex Scale infrastructure management console

You can open NetBackup from the Flex Scale infrastructure management console when you click on the **NetBackup** node.



To access NetBackup from the Flex Scale infrastructure management console

- 1** In a web browser, enter the URL for the Flex Scale infrastructure management console.

`https://ManagementServerIPorFQDN:14161/`

The *ManagementServerIP* is the public IP address or FQDN that you specified for the NetBackup Flex Scale management server.


- 2** Enter the credentials for a user with the Appliance administrator role and click **Sign in**.

- 3** On the left, click **NetBackup**.

This action launches the Flex Scale web UI in the same browser tab, where you can manage both NetBackup and Flex Scale.

Manage NetBackup and the NetBackup Flex Scale cluster management from the NetBackup Flex Scale web UI

You can manage both NetBackup and the NetBackup Flex Scale cluster management from the NetBackup Flex Scale web UI.




To access NetBackup and Flex Scale cluster management from the NetBackup Flex Scale web UI

- 1 In a web browser, enter the URL for the NetBackup Flex Scale web UI.
`https://ManagementServerIPorFQDN/webui`
The *ManagementServerIPorFQDN* is the host name or IP address of the NetBackup Flex Scale server that you want to sign in to.
- 2 Enter the credentials for a user with the Appliance Administrator role and click **Sign in**.
The web UI displays the NetBackup functionality and the NetBackup Flex Scale **Cluster management** node.

Access NetBackup Flex Scale from the NetBackup web UI

You can open NetBackup Flex Scale from the NetBackup web UI when you click on the **Appliance management** node.



To access Flex Scale from the NetBackup web UI

- 1 In a web browser, enter the URL for the NetBackup web UI.

<https://primaryserver/webui/login>

The primary server is the host name or IP address of the NetBackup primary server that you want to sign in to.

See “Sign in to the NetBackup web UI” on page 32.

- 2 Enter the credentials for a user with the Appliance administrator role and click **Sign in**.
- 3 On the left, click **Appliance management**.

In a new browser window, the NetBackup Flex Scale infrastructure management console opens.

NetBackup workloads

This chapter includes the following topics:

- [Protection of other asset types and clients](#)

Protection of other asset types and clients

The NetBackup web UI protects assets like databases, virtual machines, and clients through either protection plans or policies. Some workloads support both protection plans and policies. For more information on performing backups and restores, refer to the associated guide for that workload or agent. Protection of **Standard** and **MS-Windows** clients are covered in the *NetBackup Administrator's Guide, Volume I*.

Administering NetBackup

- [Chapter 44. Management topics](#)
- [Chapter 45. Managing client backups and restores](#)

Management topics

This chapter includes the following topics:

- [Configuring the NetBackup Client Service](#)
- [Units of measure used with NetBackup](#)
- [NetBackup naming conventions](#)
- [Wildcard use in NetBackup](#)

Configuring the NetBackup Client Service

By default, the NetBackup Client Service is configured on Windows with the **Local System** account. The **Local System** account lacks sufficient rights to perform certain backup and restore operations.

For example, for NetBackup to access CIFS volumes, the account must be changed from **Local System** to an account with access to the CIFS share.

To change the NetBackup Client Service logon account on a Windows computer:

- Open the Windows Services application.
- To change the logon account, stop the NetBackup Client Service.
- Open the properties for the NetBackup Client Service.
- Provide the name and password of the account that has the necessary permissions. For example, change the logon to that of *Administrator*.
- Restart the service.

If the logon property is not changed for the NetBackup Client Service, the policy validation fails with status code 4206.

Situations in which the NetBackup Client Service logon account requires changing

The following list contains situations in which the NetBackup Client Service logon account needs to be changed:

- To access CIFS storage for a storage unit.
- To use UNC paths, the network drives must be available to the service account that the NetBackup Client Service logs into at startup. You must change this account on each Windows client that is backed up that contains data that is shared with another computer.
- During a snapshot: To have read access to the share for backup purposes and write access during restores.
The account must be for a domain user that is allowed to access and write to the share. To verify the account, log on as that user and try to access the UNC path. For example: `\\server_name\share_name`.
- For database agents and options, configure the service with a logon account that has the necessary permission or privileges. See the documentation for your agent or option for more information.
- For the database agents that support VMware backups on a NetApp disk array, configure the logon account to one that has access to the disk array.

Units of measure used with NetBackup

For most units of measure for data, NetBackup uses the terms and abbreviations kilobyte (KB), megabyte (MB), and so on to mean the binary, or bitwise, values of each term. NetBackup does not use the powers-of-ten values, such as 1,000 for KB or 1,000,000 for MB.

When you calculate values that appear in NetBackup displays and reports, it is important to understand the difference between a unit's binary value and its powers-of-ten value. For example, a displayed value of 1.5TB actually means 1,649,267,441,664, bytes (the binary value) and not 1,500,000,000,000 bytes (the powers-of-ten value), a difference of almost 150 billion bytes.

The following table shows a number of common displayed units of measure with their corresponding bitwise names, binary multipliers, and actual values.

Table 44-1 Units of measure used in NetBackup

Displayed unit	Bitwise unit	Binary multiplier	Actual value in bytes
Kilobyte (KB)	Kebibyte (KiB)	2 ¹⁰	1024

Table 44-1 Units of measure used in NetBackup (*continued*)

Displayed unit	Bitwise unit	Binary multiplier	Actual value in bytes
Megabyte (MB)	Mebibyte (MiB)	2 ²⁰	1048576
Gigabyte (GB)	Gibibyte (GiB)	2 ³⁰	1073741824
Terabyte (TB)	Tibibyte (TiB)	2 ⁴⁰	1099511627776
Petabyte (PB)	Pebibyte (PiB)	2 ⁵⁰	1125899906842624
Exabyte (EB)	Exbibyte (EiB)	2 ⁶⁰	1152921504606846976

The Institute of Electrical and Electronics Engineers (IEEE) and the International Electrotechnical Commission (IEC) have adopted standards for these values. See the following articles for more information:

- <https://standards.ieee.org/standard/1541-2002.html> (with a paid IEEE subscription)
https://en.wikipedia.org/wiki/IEEE_1541-2002
- https://en.wikipedia.org/wiki/ISO/IEC_80000

NetBackup naming conventions

NetBackup has rules for naming logical constructs, such as clients, disk pools, backup policies, storage lifecycle policies, and so on. Generally, names are case-sensitive. The following set of characters can be used in user-defined names and passwords:

- Alphabetic (A-Z a-z) (names are case-sensitive)
- Numeric (0-9)
- Period (.)
Do not use periods in the WORM volume names.
- Plus (+)
- Hyphen (-)
Do not use a hyphen as the first character.
- Underscore (_)

These characters are also used for foreign languages.

Note: No spaces are allowed.

The Logical Storage Unit (LSU) name or the Domain Volume name must have fewer than 50 ASCII characters including a hyphen (-) and an underscore (_) and must not have a blank space.

Wildcard use in NetBackup

NetBackup recognizes the following wildcard characters in areas where wildcards can be used. (For example, in the paths of include and exclude file lists.)

The following table shows the wildcards that can be used in various NetBackup dialog boxes and lists.

Table 44-2 Wildcard use in NetBackup

Wildcard	Use
*	An asterisk serves as a wildcard for zero or more characters. An asterisk can be used in the backup selection list, the include list, and the exclude list for Windows and UNIX clients. For example: <code>r*</code> refers to all files that begin with <code>r</code> <code>r*.doc</code> refers to all files that begin with <code>r</code> and end with <code>.doc</code>. To back up all files that end in <code>.conf</code>, specify: <code>/etc/*.conf</code>
?	A question mark serves as a wildcard for any single character (A through Z; 0 through 9). A question mark can be used in the backup selection list, the include list, and the exclude list for Windows and UNIX clients. For example: <code>file?</code> refers to <code>file2</code>, <code>file3</code>, <code>file4</code> <code>file??</code> refers to <code>file12</code>, <code>file28</code>, <code>file89</code> To back up all files named <code>log01_03</code>, <code>log02_03</code>, specify: <code>c:\system\log??_03</code>

Table 44-2 Wildcard use in NetBackup (*continued*)

Wildcard	Use
[]	A pair of square brackets indicates any single character or range of characters that are separated with a dash. For example: <code>file[2-4]</code> refers to <code>file2</code>, <code>file3</code>, and <code>file4</code> <code>file[24]</code> refers to <code>file2</code>, <code>file4</code> <code>*[2-4]</code> refers to <code>file2</code>, <code>file3</code>, <code>file4</code>, <code>name2</code>, <code>name3</code>, <code>name4</code> Brackets are not valid wildcards under all circumstances for all clients: ■ Brackets that are used as wildcards in include and exclude lists: Windows clients: Allowed UNIX clients: Allowed ■ Brackets that are used as wildcards in policy backup selections lists: Windows clients: Not allowed; the use of brackets in policy backup selections lists causes backups to fail with a status 71. UNIX clients: Allowed
{ }	Curly brackets can be used in the backup selection list, the include list, and the exclude list for UNIX clients only. A pair of curly brackets (or braces) indicates multiple file name patterns. Separate the patterns by commas only; no spaces are permitted. A match is made for any or all entries. For example: <code>{*1.doc,*}.pdf</code> refers to <code>file1.doc</code>, <code>file1.pdf</code>, <code>file2.pdf</code> Note: Curly brackets are valid characters for Windows file names and cannot be used as wildcards on Windows platforms. Backslashes cannot be used as escape characters for curly bracket characters.

To use wildcard characters literally, precede the character with a backslash (\).

A backslash (\) acts as an escape character only when it precedes a special or a wildcard character. NetBackup normally interprets a backslash literally because a backslash is a legal character to use in paths.

Assume that the brackets in the following examples are to be used literally:

`C:\abc\fun[ny]name`

In the exclude list, precede the brackets with a backslash:

`C:\abc\fun\[ny\]name`

Table 44-3 Placement of wildcards in the path of backup selections

Client type	Examples
For Windows clients, wildcards function correctly only when they are placed at the end of the path, in the file or the directory name.	The following example is allowed: <pre>C:\abc\xyz\r*.doc</pre> Wildcard characters do not work elsewhere in the path. For example, an asterisk functions as a literal character (not as a wildcard) in the following examples: <pre>C:*\xyz\myfile</pre> <pre>C:\abc*\myfile</pre>
For UNIX clients, wildcards can appear anywhere in the path.	The following examples are allowed: <pre>/etc/*/abc/myfile</pre> <pre>/etc/misc/*/myfile</pre> <pre>/etc/misc/abc/*.*</pre>

Managing client backups and restores

This chapter includes the following topics:

- [About server-directed restores](#)
- [About client-redirected restores](#)
- [About restoring the files that have Access Control Lists \(ACLs\)](#)
- [About setting the original atime for files during restores on UNIX](#)
- [Restoring the System State](#)
- [About the backup and restore of compressed files on VxFS file systems](#)
- [About backups and restores on ReFS](#)

About server-directed restores

A NetBackup user with the Administrator role or similar permissions can perform restores from the NetBackup primary server. This type of restore is available in the web UI for the following policy types:

BigData	Hypervisor – Nutanix	NDMP
Cloud-Object-Store	Lotus-Notes	Nutanix-AHV
Datastore		Standard
FlashBackup	MSDP-Object-Store	Universal-Share
FlashBackup-Windows	NAS-Data-Protection	VMware (agent-based recovery)

Hyper-V

NBU-Catalog

Restore types in addition to “Normal backups” are available for certain policy types.

Restore type	Supported policy typed
Archived backups	MS-Windows, Standard
Optimized backups	MS-Windows
Point-in-time rollback	MS-Windows, NAS-Data-Protection, Standard
Raw partition backups	FlashBackup, FlashBackup-Windows, Standard
True image backups	MS-Windows, NAS-Data-Protection, NBU-Catalog, Standard
Virtual disk restore	VMware
Virtual machine backups	Hyper-V, Hypervisor-Nutanix, Nutanix-AHV

Preventing server-directed restores for a client

By default, NetBackup clients are configured to allow NetBackup administrators on a primary server to direct restores to any client.

To prevent server-directed restores for a client do the following:

- On Windows clients:
Open the **Backup, Archive, and Restore** interface.
Select **File > NetBackup Client Properties > General**, then clear the **Allow server-directed restores** check box.
- On UNIX clients:
Add `DISALLOW_SERVER_FILE_WRITES` to the following file on the client:

```
/usr/opensv/netbackup/bp.conf
```

Note: On UNIX systems, the redirected restores can incorrectly set UIDs or GIDs that are too long. The UIDs and GIDs of files that are restored from one platform to another may be represented with more bits on the source system than on the destination system. If the UID or the GID name in question is not common to both systems, the original UID or GID may be invalid on the destination system. In this case, the UID or GID is replaced with the UID or GID of the user that performs the restore.

Generating progress logs on UNIX

On UNIX, no progress log is produced if the `bp.conf` file of the requesting server does not contain an entry for the restoring server. Without that entry, the restoring server has no access to write the log files to the requesting server. (A progress log is an entry in the **Task Progress** tab of the **Backup, Archive, and Restore** client interface.)

Consider the following solutions:

- To produce a progress log, add the requesting server to the server list.
Log on to the requesting server. In the NetBackup web UI, open the host properties for the primary server. Then click **Servers**. Add the restoring server to the server list.
- Log on to the restoring server. Go to the Activity monitor to determine the success of the restore operation.

To restore a UNIX backup that contains soft and hard links, run the **Backup, Archive, and Restore** client interface from a UNIX machine.

About client-redirected restores

The **Backup, Archive, and Restore** client interface contains options for allowing clients to restore the files that were backed up by other clients. The operation is called a redirected restore.

For the following Backup Services API (XBSA) agents, redirected restores to a different version of the agent is not supported:

- MariaDB
- MySQL
- PostgreSQL

If you use a non-root service user account, specific access must be allowed for that user when you add files to the `/usr/opensv/netbackup/db/altnames` directory. The service user account must have full access to these files through the ownership or group and the permissions. For example, if the service user is `svcname` and its group is `svrgrp`, the file can have permissions of `400`. If the file owner is for a different user and group, the file permissions must allow access to the service user. For example, `777`. Equivalent permission settings must be used in a Windows environment.

About restore restrictions

By default, NetBackup permits only the client that backs up files to restore those files. NetBackup ensures that the client name of the requesting client matches the peer name that was used to connect to the NetBackup server.

Unless clients share an IP address, the peer name is equivalent to the client's host name. (Clients can share an IP address due to the use of a gateway and token ring combination, or multiple connections.) When a client connects through a gateway, the gateway can use its own peer name to make the connection.

The NetBackup client name is normally the client's short host name, such as `client1` rather than a longer form such as `client1.null.com`.

The client name is found in the following location:

Open the **File > Backup, Archive, and Restore** interface. Click **File > Specify NetBackup Machines and Policy Type**. The client name that is selected as **Source client for restores** is the source of the backups to be restored.

Allowing all clients to perform redirected restores

The NetBackup administrator can allow clients to perform redirected restores. With this, all clients can restore the backups that belong to other clients.

To do so, first create an `altnames` directory on the NetBackup primary server where the backup policy for the clients resides. Place an empty `No.Restrictions` file inside of the directory.

- On Windows:

```
install_path\NetBackup\db\altnames\No.Restrictions
```

Do not add a suffix to the files in the `altnames` directory.

- On UNIX:

```
/usr/opensv/netbackup/db/altnames/No.Restrictions
```

The NetBackup client name setting on the requesting client must match the name of the client for which the backup was created. The peer name of the requesting client does not need to match the NetBackup client name setting.

Note: The `altnames` directory can present a breach of security, so use it only under limited circumstances. Users that are permitted to restore files from other clients may also have local permission to create the files that are found in the backup.

Caution: For security reasons, it is strongly advised not to use the `No.Restrictions` file approach. This approach enables any client to restore backups of other clients that can be a security threat.

Note: On using the `No.Restrictions` file approach, a notification is by default generated in the NetBackup web UI every 7 days. Use the `NOTIFY_SNOOZE_PERIOD_IN_DAYS` option to change the frequency of this notification from the default value to any value from 1 to 90.

For information on alternative methods of alternate client restores, refer to the following topics:

See [“Allowing a single client to perform redirected restores”](#) on page 622.

See [“Allowing redirected restores of a specific client’s files”](#) on page 622.

Allowing a single client to perform redirected restores

The NetBackup administrator can permit a single client to restore the backups that belong to other clients.

To do so, create an `altnames` directory on the NetBackup primary server where the policy that backed up the other client(s) resides. Place an empty *peername* file inside of the `altnames` directory where *peername* is the client to possess restore privileges.

- On Windows:

```
install_path\NetBackup\db\altnames\peername
```

- On UNIX:

```
/usr/opensv/netbackup/db/altnames/peername
```

In this case, the requesting client (*peername*) can access the files that are backed up by another client. The NetBackup client name setting on *peername* must match the name of the other client.

Allowing redirected restores of a specific client’s files

The NetBackup administrator can permit a single client to restore the backups that belong to another specific client.

To do so, create an `altnames` directory on the NetBackup primary server of the requesting client in the following location:

- On Windows:

```
install_path\NetBackup\db\altnames\peername
```

- On UNIX:

`/usr/opensv/netbackup/db/altnames/peername`

Then, create a *peername* file inside of the directory where *peername* is the client to possess restore privileges. Add to the *peername* file the names of the client(s) whose files the requesting client wants to restore.

The requesting client can restore the files that were backed up by another client if:

- The names of the other clients appear in the *peername* file, and
- The NetBackup client name of the requesting client is changed to match the name of the client whose files the requesting client wants to restore.

Examples of redirected restores

This topic provides some example configurations that allow clients to restore the files that were backed up by other clients. These methods may be required when a client connects through a gateway or has multiple Ethernet connections.

In all cases, the requesting client must have access to an image database directory on the primary server or the requesting client must be a member of an existing NetBackup policy.

- On Windows: `install_path\NetBackup\db\images\client_name`
- On UNIX: `/usr/opensv/netbackup/db/images/client_name`

Note: Not all file system types on all computers support the same features. Problems can be encountered when a file is restored from one file system type to another. For example, the S51K file system on an SCO computer does not support symbolic links nor does it support names greater than 14 characters long. You may want to restore a file to a computer that doesn't support all the features of the computer from which the restore was performed. In this case, all files may not be recovered.

In the following examples, assume the following conditions:

- *client1* is the client that requests the restore.
- *client2* is the client that created the backups that the requesting client wants to restore.
- On Windows: `install_path` is the path where you installed the NetBackup software. By default, this path is `C:\Program Files\Veritas`.

Note: The information in this topic applies to the restores that are made by using the command line, not the **Backup, Archive, and Restore** client interface.

Note: On Windows: You must have the necessary permissions to perform the following steps.

On UNIX: You must be a root user for any of the steps that must be performed on the NetBackup server. You may also need to be a root user to make the changes on the client.

Example of a redirected client restore

Assume you must restore files to *client1* that were backed up from *client2*. The *client1* and *client2* names are those specified by the NetBackup client name setting on the clients.

On Windows:

- 1 Log on to the NetBackup server.
- 2 Add *client2* to the following file and perform one of the following:
 - Edit `install_path\NetBackup\db\altnames\client1` to include the name of *client2*.
 - Create the following empty file:
`install_path\NetBackup\db\altnames\No.Restrictions`

On UNIX:

- 1 Log on as root on the NetBackup server.
- 2 Perform one of the following actions:
 - Edit `/usr/opensv/netbackup/db/altnames/client1` so it includes the name of *client2*. Or,
 - Run the `touch` command on the following file:
`/usr/opensv/netbackup/db/altnames/No.Restrictions`

Note: The `No.Restrictions` file allows any client to restore files from *client2*.

- 3 Log on to *client1* and change the NetBackup client name to *client2*.
- 4 Restore the file.
- 5 Undo the changes that were made on the server and client.


Example of a redirected client restore using the altnames file

This example explains how `altnames` provides restore capabilities to clients that do not use their own host name when they connect to the NetBackup server.

By default, the NetBackup client name of the requesting client must match the peer name that is used in the connection to the NetBackup server. When the NetBackup client name is the host name for the client and matches the peer name (normal case), this requirement is met.

However, problems arise when clients connect to multiple ethernet or connect to the NetBackup server through a gateway.

Figure 45-1 Example restore from a token ring client



In this example, restore requests from *client1*, *client2*, and *client3* are routed through the TCP gateway. Because the gateway uses its own peer name rather than the client host names for connection to the NetBackup server, NetBackup refuses the requests. Clients cannot restore even their own files.

To correct the situation, do the following

- 1 Determine the peer name of the gateway:
 - Try a restore from the client in question. In this example, the request fails with an error message similar to the following:
`client is not validated to use the server`
 - Examine the NetBackup problems report and identify the peer name that is used on the request. Entries in the report may be similar to the following:
`01/29/12 08:25:03 bpserver - request from invalid server or
client client1.dvlp.null.com`

In this example, the peer name is `client1.dvlp.null.com`.

2 Do one of the following:

On Windows: Determine the peer name, then create the following file on the NetBackup primary server:

```
install_path\NetBackup\db\altnames\peername
```

In this example, the file is:

```
install_path\NetBackup\db\altnames\client1.dvlp.null.com
```

On UNIX: Run the `touch` command on the following file:

```
/usr/openv/netbackup/db/altnames/peername
```

In this example, the file is:

```
/usr/openv/netbackup/db/altnames/client1.dvlp.null.com
```

3 Edit the *peername* file so that it includes the client names.

For example, if you leave file `client1.dvlp.null.com` empty, *client1*, *client2*, and *client3* can all access the backups that correspond to their NetBackup client name setting.

See [“Allowing a single client to perform redirected restores”](#) on page 622.

If you add the names *client2* and *client3* to the file, you give these two clients access to NetBackup file restores, but exclude *client1*.

See [“Allowing redirected restores of a specific client’s files”](#) on page 622.

Note that this example requires no changes on the clients.

4 Restore the files.

Example of how to troubleshoot a redirected client restore using the altnames file

If you cannot restore files with a redirected client restore by using the `altnames` file, troubleshoot the situation, as follows.

On Windows:

1 Create the debug log directory for the NetBackup Request Daemon:

```
install_path\NetBackup\logs\bprd
```

2 On the primary server, stop and restart the NetBackup Request Daemon. Restart the service to ensure that this service is running in verbose mode and logs information regarding client requests.

3 On *client1* (the requesting client), try the file restore.

- 4 On the primary server, identify the peer name connection that *client1* uses.
- 5 Examine the debug log for the NetBackup Request Daemon to identify the failing name combination:

Examine the failure as logged on the **All log entries** report. Or, examine the debug log for the NetBackup Request Daemon to identify the failing name combination:

```
install_path\NetBackup\logs\bprd\mmddyy.log
```

- 6 On the primary server, do one of the following:
 - Create an *install_path*\NetBackup\db\altnames\No.Restrictions file. The file allows any client to access *client2* backups if the client changes its NetBackup client name setting to *client2*.
 - Create an *install_path*\NetBackup\db\altnames\peername file. The file allows *client1* to access *client2* backups if *client1* changes its NetBackup client name setting to *client2*.
 - Add *client2* name to the following file:

```
install_path\NetBackup\db\altnames\peername.
```
 - *client1* is allowed to access backups on *client2* only.
- 7 On *client1*, change the NetBackup client name setting to match what is specified on *client2*.
- 8 Restore the files from *client1*.
- 9 Perform the following actions:
 - Delete *install_path*\NetBackup\logs\bprd and the contents.
 - In the NetBackup web UI, open the host properties for the primary server. Click **Logging**. Clear the **Keep logs for days** setting.
- 10 If you do not want the change to be permanent, do the following:
 - Delete *install_path*\NetBackup\db\altnames\No.Restrictions (if existent).
 - Delete *install_path*\NetBackup\db\altnames\peername (if existent).
 - On *client1*, change the NetBackup client name to its original value.

On UNIX:

- 1 On the NetBackup primary server, add the `VERBOSE` entry and a logging level to the `bp.conf` file. For example:

```
VERBOSE = 3
```

- 2 Create the debug log directory for `bprd` by running the following command:

```
mkdir /usr/opensv/netbackup/logs/bprd
```

- 3 On the NetBackup server, stop the NetBackup Request Daemon, `bprd`, and restart it in verbose mode by running:

```
/usr/opensv/netbackup/bin/admincmd/bprdreq -terminate  
/usr/opensv/netbackup/bin/bprd -verbose
```

Restart `bprd` to ensure that `bprd` logs information regarding client requests.

- 4 On *client1*, try the file restore.
- 5 On the NetBackup server, identify the peer name connection that *client1* used.

Examine the `bard debug` log to identify the failing name combination:

Examine the failure as logged on the **All log entries** report or examine the `bard debug` log to identify the failing name combination:

```
/usr/opensv/netbackup/logs/bprd/log.date
```

- 6 On the NetBackup server enter the following command:

```
mkdir -p /usr/opensv/netbackup/db/altnames touch  
/usr/opensv/netbackup/db/altnames/No.Restrictions
```

This command allows any client access to *client2* backups by changing its NetBackup client name setting to specify the *client2*.

- 7 Run the touch command on the following file:

```
/usr/opensv/netbackup/db/altnames/peername
```

The command allows *client1* access to any *client2* backups by changing its NetBackup client name setting to specify *client2*.

- 8 Add *client2* to the `/usr/opensv/netbackup/db/altnames/peername` file. The addition to the `peername` file allows *client1* access to the backups that were created on *client2* only.
- 9 On *client1*, change the NetBackup client name setting in the user interface to match what is specified on *client2*.
- 10 Restore the files to *client1*.

11 Do the following:

- Delete the `VERBOSE` entry from the `/usr/opensv/netbackup/bp.conf` file on the primary server.
- Delete `/usr/opensv/netbackup/logs/bprd` and the contents.

12 Return the configuration to what it was before the restore.

- Delete `/usr/opensv/netbackup/db/altnames/peer.or.hostname` (if it exists)
- Delete `/usr/opensv/netbackup/db/altnames/No.Restrictions` (if it exists)
- On *client1*, restore the NetBackup client name setting to its original value.

About restoring the files that have Access Control Lists (ACLs)

An Access Control List (ACL) is a table that conveys the access rights users need to a file or directory. Each file or directory can have a security attribute that extends or restricts users' access.

By default, the `nbtar` (`/usr/opensv/netbackup/bin/nbtar`) restores ACLs along with file and directory data.

However, in some situations the ACLs cannot be restored to the file data, as follows:

- Where the restore is cross-platform.
- When a restore utility (`tar`) other than `nbtar` is used to restore files.

In these instances, NetBackup stores the ACL information in a series of generated files in the `root` directory using the following naming form:

`.SeCuRiT.y.nnnn`

These files can be deleted or can be read and the ACLs regenerated by hand.

Note: If performing an alternate restore where the original directory was ACL-enabled, the alternate restore directory must also be ACL-enabled. If the alternate restore directory is not ACL-enabled, the restore is not successful.

Restoring files without restoring ACLs

The NetBackup client interface on Windows is available to administrators to restore data without restoring the ACLs. Both the destination client and the source of the backup must be Windows systems.

To restore files without restoring ACLs, the following conditions must be met:

- The policy that backed up the client is of policy type **MS-Windows**.
- An administrator performs the restore and is logged into a NetBackup server (Windows or UNIX). The option is set at the server by using the client interface. The option is unavailable on standalone clients (clients that do not contain the NetBackup server software).
- The destination client and the source of the backup must both be systems running supported Windows OS levels. The option is disabled on UNIX clients.

Use the following procedure to restore files without restoring ACLs.

To restore files without restoring ACLs

- 1 Log on to the NetBackup server as administrator.
- 2 Open the **Backup, Archive, and Restore** client interface.
- 3 From the client interface, initiate a restore.
- 4 Select the files to be restored, then select **Actions > Start Restore of Marked Files**.
- 5 In the **Restore Marked Files** dialog box, place a check in the **Restore without access-control attributes** check box.
- 6 Make any other selections for the restore job.
- 7 Click **Start Restore**.

About setting the original atime for files during restores on UNIX

During a restore, NetBackup sets the `atime` for each file to the current time by default. You can elect to have NetBackup set the `atime` for each restored file to the value the file had when it was backed up. To do so, create the following file on the client:

```
/usr/opensv/netbackup/RESTORE_ORIGINAL_ETIME
```

Restoring the System State

The System State includes the registry, the COM+ Class Registration database, and boot and system files. If the server is a domain controller, the data also includes the Active Directory services database and the SYSVOL directory.

Note: The best recovery procedure depends on many hardware and software variables that pertain to the server and its environment. For a complete Windows recovery procedure, refer to the Microsoft documentation.

Read the following notes carefully before you restore the System State:

- The System State should be restored in its entirety. Do not restore selected files.
- Do not redirect a System State restore. System State is computer-specific and to restore it to an alternate computer can result in an unusable system.
- Do not cancel a System State restore operation. To cancel the operation may leave the system unusable.
- To restore the System State to a domain controller, the Active Directory must not be running.

Restoring the System State

Use the following procedure to restore the System State.

To restore the System State

- 1 To restore the Active Directory, restart the system, and press F8 during the boot process. F8 brings up a startup options menu. Press F8 upon restart if the system to which you are to restore is a Windows domain controller. Otherwise, begin with step 4.
- 2 From the startup options, select **Directory Services Restore Mode** and continue the boot process.
- 3 Ensure that the **NetBackup Client Service**, either `bpinetd` on Windows or `inetd` on UNIX, has started. Use the **Activity Monitor** or the Services application in the Windows Control Panel.
- 4 Start the **Backup, Archive, and Restore** client interface. Click **Select for Restore**, and place a checkmark next to **System State**.
- 5 To restore a system state backup using an incremental backup, select the full backup and one or more differential-incremental or cumulative-incremental backups.
- 6 From the **Actions** menu, select **Restores**.
- 7 From the **Restore Marked Files** dialog box, select **Restore everything to its original location** and **Overwrite the existing file**.

Do not redirect the System State restore to a different host. System State is computer-specific. To restore it to a different computer can result in an unusable system.

- 8 Click **Start Restore**.
- 9 The network may contain more than one domain controller. To replicate Active Directory to other domain controllers, perform an authoritative restore of the Active Directory after the NetBackup restore job completes.

To perform an authoritative restore of the Active Directory, run the Microsoft `ntdsutil` utility after you restored the System State data but before the server is restarted. An authoritative restore ensures that the data is replicated to all of the servers.

Additional information about an authoritative restore and the `ntdsutil` utility is available.

See the Microsoft documentation.
- 10 Restart the system before performing subsequent restore operations.

If you booted into **Directory Services Restore Mode** on a domain controller, restart into normal mode when the restore is complete.

About the backup and restore of compressed files on VxFS file systems

NetBackup can back up and restore VxFS-compressed files, maintaining the compression state when the target volume supports file system compression. Future releases will expand this capability to other file systems.

Upon backup of files on a VxFS file system, a message displays in the **Activity Monitor** whenever NetBackup encounters a compressed file:

```
Compress flag found for 'file_name'.
```

Upon restore, NetBackup restores the files to a VxFS file system in their compressed form.

If the restore is to a non-VxFS file system, NetBackup restores the files in an uncompressed form. The following message displays in the **Progress** tab of the **Backup, Archive, and Restore** client interface:

```
File 'file_name' will not be restored in compressed form. Please refer to the Release Notes or User Guide.
```

The message appears only for the first file that cannot be restored in its compressed form.

Note: The compression messages display if the verbose level is 1 or greater.

About backups and restores on ReFS

Microsoft Resilient File System (ReFS) support in NetBackup is automatic and requires no additional configuration.

NetBackup does not support a redirected restore of a Microsoft Resilient File Systems (ReFS) file system.

[Table 45-1](#) lists the ReFS-to-NTFS backup and restore combinations and the success of each.

Table 45-1 ReFS backup and restore

Between file systems	Backups	Restores
ReFS to ReFS	Successful	Successful
ReFS to NTFS	Successful	Successful
NTFS to ReFS	Successful	Limited success For successful restores: ■ Restore NTFS backups to NTFS file system. ■ Remove all non-supported ReFS items. ■ Copy the files to an ReFS file system.

Known issue

A known issue exists that includes failures with respect to backups for files having ReFS based snapshot. At present Microsoft does not support backup of files having ReFS based snapshot as the API's are not compatible. Microsoft is working on documenting this behavior and providing support which are tracked with the following issue ID's:

- Documentation issue#: 42324557
- Backup Read issue#: 42295538

Disaster recovery and troubleshooting

- [Chapter 46. Disaster recovery of NetBackup](#)
- [Chapter 47. Managing Resiliency Platforms](#)
- [Chapter 48. Managing Bare Metal Restore \(BMR\)](#)
- [Chapter 49. Troubleshooting the NetBackup Web UI](#)

Disaster recovery of NetBackup

This chapter includes the following topics:

- [About disaster recovery of NetBackup](#)

About disaster recovery of NetBackup

Disaster recovery of NetBackup is discussed in the [NetBackup Troubleshooting Guide](#).

That guide includes the following types of information:

- Disk recovery procedures
- Clustered NetBackup server recovery
- Restoring the disaster recovery package
- Recovering the NetBackup catalog

Managing Resiliency Platforms

This chapter includes the following topics:

- [About Resiliency Platform in NetBackup](#)
- [Understanding the terms](#)
- [Configuring a Resiliency Platform](#)
- [Troubleshooting NetBackup and Resiliency Platform issues](#)

About Resiliency Platform in NetBackup

You can integrate NetBackup and Veritas Resiliency Platform to manage your disaster recovery operations. Veritas Resiliency Platform provides a single console from which you can proactively maintain business uptime across private, public, and hybrid clouds. Integrating NetBackup and Resiliency Platform lets you leverage the capabilities, such as complete automation, visualizing and monitoring DR specific information for all resiliency operations for the virtual machines in your data center.

Note the following points:

- You can integrate more than one Resiliency Platform with your NetBackup primary server.
- You can have more than one data centers for a Resiliency Platform.
- You can use Resiliency Platform with Veritas Resiliency Platform version 3.5 and later in NetBackup.
- After you add a Resiliency Platform, the assets are automatically discovered and displayed on the **Virtual machines** tab.

- You can view detailed information alerts and error messages in the **Notifications** section.

Understanding the terms

The following table explains the key components related to Veritas Resiliency Platform and NetBackup integration.

Term	Description
Resiliency Platform	The Veritas Resiliency Platform integrated with your NetBackup primary server. The Resiliency Manager provides the services required for protecting assets, such as virtual machines, within a resiliency domain. It also provides workload automation services.
Resiliency manager	The component that provides resiliency capabilities within a resiliency domain. It is composed of loosely coupled services, a distributed data repository, and a management console.
Infrastructure management server (IMS)	The component that discovers, monitors, and manages the asset infrastructure within a data center. The IMS transmits information about the asset infrastructure to the Resiliency Manager. The IMS is deployed as a virtual appliance. To achieve scale, multiple IMSs can be deployed in the same data center.
Data center	The location that contains source data center and a target data center. Each data center has one or more IMSs.
Resiliency group	The unit of management and control in Resiliency Platform. You organize related assets into a resiliency group and manage and monitor them as a single entity.
Automated virtual machines	The assets that are a part of a resiliency group and you can perform actions, such as migrate, recover, and rehearsal.
Recovery readiness	Measured based on migrate, recover or rehearsal operations. ■ Low - If no operations are performed or failed. ■ High - If at least one operation is performed successfully in the past 7 days. ■ Medium - If the recovery readiness does not fall in either high or low category.

Term	Description
Recovery Point Object (RPO)	Recovery Point Objective is the point in time you can recover to in the event of a disaster. For example, if you have an RPO of 4 hours on your critical virtual machines then you lose 4 hours of data, as 4 hours ago is the last point in time to which you can recover data on your VMs.

Configuring a Resiliency Platform

You can add, edit, delete, or refresh a Resiliency Platform. You can add more than one Resiliency Platform in NetBackup.

Add a Resiliency Platform

You can add one or more than one Resiliency Platforms in NetBackup. The Resiliency Platform lets you add virtual machines and automate protection. If the resiliency manager is using a third-party certificate, see the [NetBackup Web UI Administrator's Guide](#).

To add a Resiliency Platform

- 1 On the left, click **Resiliency**.
- 2 Click the **Resiliency Platform** tab.
- 3 Click **Add Resiliency Platform**.
- 4 Read the instructions on the **Add Resiliency Platform** dialog box and click **Next**.
- 5 In the **Add credentials** dialog box, enter a value in the following fields and click **Next**:
 - **Resiliency manager host name or IP address**
 - **Resiliency Platform API access key**
 - **NetBackup API access key**
- 6 In the **Add data center and Infrastructure management** server dialog box, select a data center.
- 7 In the **Infrastructure management server** section, select a preferred server.
- 8 Click **Add**.

After you add the Resiliency Platform in NetBackup, the NetBackup primary server will be configured automatically in the Resiliency Platform.

Note: If the NetBackup has FIPS mode enabled and you need to fetch the respective certificates, refer *Integrating with NetBackup* topic in *Resiliency Platform* product documentation. You need to install Resiliency Platform certificates in FIPS trust store and then add the Resiliency Platform. (Only done when NetBackup has FIPS mode enabled)

Configure a third-party CA certificate

You can use a self-signed or a third-party certificate to validate your Resiliency manager.

Consider the following points:

- For Windows, you can give a certificate as a file path or install the third-party certificate in the Trusted root certificates authorities.
- To switch from a self-signed certificate to a third-party certificate for an already added Resiliency Platform, you can edit the Resiliency Platform.

To configure a third-party CA certificate

- 1 Copy a PKCS #7 or P7B file having certificates of the trusted root certificates authorities that are bundled together. This file may either be PEM or DER encoded.
- 2 Create a CA file containing the PEM encoded certificates of the trusted root certificate authorities that are concatenated together.
- 3 In the `bp.conf` file, create the following entries, where `/certificate.pem` is the file name:
 - `ECA_TRUST_STORE_PATH = /certificate.pem`
 - Verify that the `nbwebsvc` account has the permissions to access the path that `ECA_TRUST_STORE_PATH` refers.

Edit or delete a Resiliency Platform

After you add a Resiliency Platform, you can edit the Resiliency Platform and NetBackup API access keys. You cannot change or update the Resiliency manager host name or IP address. However, you can delete the Resiliency Platform and add it to NetBackup again. If you refresh the Resiliency Platform, the discovery of assets on the Resiliency Platform is triggered.

To edit a Resiliency Platform

- 1 On the left, click **Resiliency**.
- 2 Click the **Resiliency Platform** tab.

- 3 Click the **Actions** menu for the Resiliency Platform that you want to edit and select **Edit**.
- 4 Enter the updated **Resiliency Platform API access key** and **NetBackup API access key**.
- 5 Click **Next**.
- 6 In the **Edit data center and Infrastructure management server** dialog box, select the **Data center** and then select the preferred infrastructure management server.
- 7 Click **Save**.
- 8 To delete a Resiliency Platform, from the **Actions** menu, select **Delete**.

View the automated or not-automated VMs

The virtual machines that belong to a resiliency group in Veritas Resiliency Platform are discovered and displayed on the **Automated** tab and the VMs that don't belong any resiliency group are displayed on the **Not automated** tab. You can view the status of the assets and perform various actions. You can search for a VM or apply filters too.

The following table lists the columns displayed on the **Automated** and **Not automated** tabs:

Table 47-1

Tab	Column	Description
■ Automated ■ Not automated	Name	Name of the virtual machine.
■ Automated	RPO	Recovery Point Objective is the point in time you can recover to in the event of a disaster. For example, if you have an RPO of 4 hours on your critical virtual machines then you lose 4 hours of data, as 4 hours ago is the last point in time to which you can recover data on your VMs.
■ Automated ■ Not automated	State	Whether the VM is switched on or off.

Table 47-1 (continued)

Tab	Column	Description
Automated	Recovery readiness	Measured based on migrate, recover or rehearsal operations. - Low - If no operations are performed or failed. - High - If at least one operation is performed successfully in the past 7 days. - Medium - If the recovery readiness does not fall in either high or low category.
- Automated - Not automated	Platform	The platform that the VM belongs to.
- Automated - Not automated	Server	The server name of the VM.
Automated	Protection	Protection status of the VM.
Automated	Resiliency group	Name of the resiliency group to which the VM belongs.
Not automated	Recovery action	Launch the Resiliency Platform to add the VM to a resiliency group.

To view and perform actions on automated VMs

- 1 On the left, click **Resiliency**.
- 2 On the **Virtual machines** tab, click **Automated**.
- 3 To view more details about a VM, in the **Name** column, click a VM.
- 4 To view all VMs that are a part of the same resiliency group, click the preferred resiliency group.

- 5 To perform disaster recovery operation, such as rehearse, restore, or recover, click **Launch Resiliency Platform**.

To enable single-sign, same authentication domain must be configured NetBackup and Veritas Resiliency Platform. If not configured, you must login with username and password to access Veritas Resiliency Platform web console.
- 6 Log on to your Resiliency Platform and perform the preferred action. See the *Veritas Resiliency Platform User Guide*.

To view and perform actions on not automated VMs

- 1 On the left, click **Resiliency**.
- 2 On the **Virtual machines** tab, click **Not automated**.
- 3 To add the VM to a resiliency group, in the **Recovery action** column, click **Automate Recovery**.
- 4 Perform the preferred action for your Resiliency Platform. See the *Veritas Resiliency Platform User Guide*.

Troubleshooting NetBackup and Resiliency Platform issues

Use the following information to troubleshoot issues.

Table 47-2 Troubleshooting issues

Issue	Action
Failed to configure the current NetBackup primary server with the Resiliency Platform.	Check the logs at the following location in Cohesity Resiliency Platform's Resiliency manager: ■ /var/opt/VRTSitrp/logs/copydata-service.log ■ /var/opt/VRTSitrp/logs/api-service.log
Failed to establish a persistent connection between the current NetBackup primary server and the Resiliency Platform.	■ Verify that the logged in user has permissions in credentials namespace. ■ Check the logs at the following location on the NetBackup primary server: ■ /usr/openv/logs/nbweb service/ in NetBackup installation directory ■ C:\Program Files\Veritas\NetBackup\logs\nbweb service in NetBackup windows

Table 47-2 Troubleshooting issues (*continued*)

Issue	Action
Failed to launch the Cohesity Resiliency Platform	Verify that same authentication domain is used to configure Cohesity Resiliency Platform and NetBackup.

Managing Bare Metal Restore (BMR)

This chapter includes the following topics:

- [About Bare Metal Restore \(BMR\)](#)
- [Add a custom role for a Bare Metal Restore \(BMR\) administrator](#)

About Bare Metal Restore (BMR)

NetBackup Bare Metal Restore (BMR) is the server recovery option of NetBackup. BMR automates and streamlines the server recovery process so you do not have to reinstall the operating systems or configure the hardware manually. BMR restores the operating system, the system configuration, and all the system files and the data files with the following steps.

For complete information on BMR, refer to the [NetBackup Bare Metal Restore Administrator's Guide](#).

In the NetBackup web UI, you can perform the following BMR operations:

- View and manage the clients that are backed up for VM conversion.
- Convert BMR-enabled backups to a virtual machine using the Virtual Machine Conversion wizard.
- Create point-in-time restore configurations.
- View and manage VM conversion tasks.
- View and manage the BMR clients and configurations.
- Run pre-restore operations on the client configuration and the VM conversion client's configurations. For example, prepare-to-restore, prepare-to-discover, and dissimilar disk restore operations.

- View and manage boot servers.
- View and manage resources like shared resource trees, discovered configurations, and Windows device driver packages.
- View and manage BMR restore or discover tasks.

Add a custom role for a Bare Metal Restore (BMR) administrator

To add a custom RBAC role

- 1 On the left, select **Security > RBAC** and click **Add**.
- 2 Select **Custom role** to manually configure all the permissions for the role.
- 3 Provide a **Role name** and a description.
For example, you may want to indicate that the role is for any users that are BMR administrators.
- 4 On the **Global** tab, expand the **BMR** section and select all the permissions for **BMR**.

Boot servers	View, Delete
Clients	View, Create, Update, Delete, Pre restore
VM conversion	View, Delete, VM conversion

- 5 Expand the **NetBackup management** section.

- Locate the **NetBackup hosts** group.
- Select the following permissions:

NetBackup hosts	View, Update

- Locate the **NetBackup backup images** group.
- Select the following permissions:

NetBackup backup images	Image Requests > View
NetBackup backup images	View

6 For ESXi servers, additional permissions are needed for **Host properties**.

- On the **Global** tab, expand the **NetBackup management** section.
- Select the following permissions:

Access hosts View, Create, Update, Delete

7 On the **Assets** tab, select the following permissions.

VMware assets View, Update, View restore targets

8 Click **Assign**.

9 Under **Workloads**, click **Assign**.

Select the VMware assets that you want the role to have access to.

- To give the role access to all VMware assets and future assets that you add, select **Apply selected permissions to all existing and future VMware assets**.
- To select individual assets, deselect **Apply selected permissions to all existing and future VMware assets** and click **Add**.
 For example, you can select one or more: datastores, datastore clusters, ESXi servers, ESXi clusters, resource pools, vApps.

10 When you have added all the assets, click **Assign**.

11 On the **Users** card, click **Assign**. Then add each user that you want to have access to this custom role.

12 When you are done configuring the role, click **Save**.

Troubleshooting the NetBackup Web UI

This chapter includes the following topics:

- [Tips for accessing the NetBackup web UI](#)
- [If a user doesn't have the correct permissions or access in the NetBackup web UI](#)
- [Unable to validate the user or group when configuring LDAP server](#)

Tips for accessing the NetBackup web UI

When NetBackup is properly configured, a user can access the primary server at the following URL:

`https://primaryserver/webui/login`

If the web UI on a primary server does not display, follow these steps to troubleshoot the issue.

Browser displays an error that the connection was refused or that it cannot connect to the host

Table 49-1 Solutions when the web user interface does not display

Step	Action	Description
Step 1	Check the network connection.	
Step 2	Verify that the firewall is open for port 443.	Refer to the following article: https://www.veritas.com/docs/100042950

Table 49-1 Solutions when the web user interface does not display
(continued)

Step	Action	Description
Step 3	If port 443 is in use, configure another port for the web UI.	Refer to the following article: https://www.veritas.com/docs/100042950
Step 4	Verify that the nbweb service is up.	Check the nbweb service logs for more details.
Step 5	Verify that the vnetd -http_api_tunnel is running.	Verify that the vnetd -http_api_tunnel service is running. For more details, check the vnetd -http_api_tunnel logs with OID 491.
Step 6	Ensure that the external certificate for the NetBackup web server is accessible and has not expired.	■ Use the Java Keytool commands to validate the following file: Windows: <code>install_path\var\global\wsl\credentials\nbweb service.jks</code> UNIX: <code>/usr/opensv/var/global/wsl/credentials nbweb service.jks</code> ■ Check whether the nbweb group has a permission to access the nbweb service.jks file. ■ Contact Cohesity Technical Support.

Cannot access web UI when you use a custom port

- Restart the vnetd service.
- Follow the steps in [Table 49-1](#).

Certificate warning displays when you try to access the web UI

The certificate warning displays if the NetBackup web server uses a certificate that is issued by a CA that is not trusted by the web browser. (Including the default NetBackup web server certificate that the NetBackup CA issued.)

To resolve a certificate warning from the browser when you access the web UI

- 1 Configure the external certificate for the NetBackup web server.
See [“Configure an external certificate for the NetBackup web server”](#) on page 429.
- 2 If the problem persists, contact Cohesity Technical Support.

If a user doesn't have the correct permissions or access in the NetBackup web UI

Note that only administrators and root users automatically have full access to the web UI. Other users must be configured in RBAC to have access and permissions for the web UI.

See [“Configuring RBAC”](#) on page 517.

If a user does not have the correct permissions or cannot access the workload assets that they should have access to, do the following:

- Verify that the user's credentials match the username (or the username and the domain name) that is specified in the user's role.
- Review the roles for the user in **Security > RBAC**. You may need to change the role permissions. However, be aware that those kinds of changes also affect any other users that belong to those roles.
- Any user account changes with the identity provider are not synchronized with the user's roles. If a user account changes with the identity provider, the user may not have the correct permissions or access. The NetBackup security administrator must edit each role for the user to remove the existing user account and re-add the new account.
- Changes to a user's roles are not immediately reflected in the web UI. A user with an active session must sign out and sign in again before any changes take effect.

Unable to validate the user or group when configuring LDAP server

When the administrator configures the LDAP server, they must specify the `-d DomainName` option. `DomainName` can be the LDAP server name or the domain name. Whatever name is specified for `-d DomainName` is the domain name that an administrator should use when they add users to an RBAC role.

If you specify the incorrect domain, you may see the error `Unable to validate the user or group`. Review the following:

- The username and domain name are typed correctly.
- You specified the correct domain name.
 The domain name that you should specify depends on how the LDAP server is configured in NetBackup. Contact your administrator for help with adding users to RBAC.

Other topics

- [Chapter 50. Additional NetBackup catalog information](#)
- [Chapter 51. About the NetBackup database](#)

Additional NetBackup catalog information

This chapter includes the following topics:

- [Parts of the NetBackup catalog](#)
- [Archiving the catalog and restoring from the catalog archive](#)
- [Estimating catalog space requirements](#)
- [About the file hash search in NetBackup](#)

Parts of the NetBackup catalog

The NetBackup catalog resides on the NetBackup primary server. It manages and controls access to the following types of data:

- Image metadata (information about backup images and copies).
- Backup content data (information about the folders, files, and the objects in a backup (. ϵ files)).
- NetBackup backup policies.
- NetBackup licensing data.
- The NetBackup error log.
- The client database.
- Cloud configuration files.

See [“About the catalog backup of cloud configuration files”](#) on page 656.

The catalog consists of the following parts:

- NetBackup stores information in the NetBackup database (NBDB). The metadata includes information about the data that has been backed up, and about where the data is stored.
See [“NetBackup databases and configuration files”](#) on page 652.
- The image database.
The image database contains information about the data that has been backed up.
See [“About the NetBackup image database”](#) on page 654.
- NetBackup configuration files.
- The key management service (KMS) configuration files
For more details on the KMS configuration, see the [NetBackup Security and Encryption Guide](#).

NetBackup is sensitive to the location of the primary server components. Running any part of NetBackup (the binaries, the logs, the database, the images) on a network share (NFS, for example) can affect performance of even normal operations. NetBackup can be CIFS-mounted on SAN or NAS storage as long as the average I/O service times remain less than 20 milliseconds.

The storage must also meet certain conditions to ensure data integrity in the NetBackup catalog.

- The order of file writes must be guaranteed.
- When a write request is issued, the write must complete to the physical storage. The write request must not merely be buffered when the SAN or the NAS returns from the write call.
See the following article for more information:

NetBackup databases and configuration files

The NetBackup catalog backup includes the NetBackup databases and the configuration files, as follows.

Databases

The NetBackup databases include the NBDB database and the NetBackup Authorization database (NBAZDB). If Bare Metal Restore is installed (optionally-licensed) there is also the BMRDB database.

The databases are located in the following directories:

`install_path\NetBackupDB\data`

`/usr/opensv/db/data/`

These directories contain the following subdirectories:

`\bmrdb\` or `/bmrdb/` (if BMR is installed)

`\nbazdb\` or `/nbazdb/` (NetBackup authorization)

`\nbdb\` or `/nbdb/` (contains both the NBDB and the EMM databases)

Configuration files

Warning: Do not edit the configuration files. NetBackup may not start if you change these files.

Note: The catalog backup process copies this data to `/usr/opensv/db/staging` and backs up the copy.

The following configuration files are created:

```
pgbouncer.ini
pg_hba.conf
pg_ident.conf
postgresql.auto.conf
postgresql.conf
userlist.txt
vxdbs.conf
web.conf
```

Most of the configuration files are located in the following directories:

```
install_path\NetBackupDB\data\instance
/usr/opensv/db/data/instance
```

`web.conf` is created in the following directories:

```
/usr/opensv/var/global/wsl/config
install_path\NetBackup\var\global\wsl\config
```

About the Enterprise Media Manager (EMM)

The Enterprise Media Manager (EMM) is a NetBackup service that manages the device and the media information for NetBackup. The Enterprise Media Manager stores its managed information in a database that resides on the primary server.

The NetBackup Resource Broker queries EMM to allocate storage units, drives (including drive paths), and media.

EMM contains the following information:

- Device attributes
- Robotic library and standalone drive residence attributes
- NDMP attributes
- Barcode rule attributes
- Volume pool attributes
- Tape attributes
- Media attributes
- Storage unit attributes
- Storage unit group attributes
- Hosts with assigned tape drives
- Media and device errors
- Disk pool and disk volume attributes
- Storage server attributes
- Log on credentials for storage servers, disk arrays, and NDMP hosts
- Fibre Transport attributes

EMM ensures consistency between drives, robotic libraries, storage units, media, and volume pools across multiple servers. EMM contains information for all media servers that share devices in a multiple server configuration. The NetBackup scheduling components use EMM information to select the server, drive path, and media for jobs.

About the NetBackup image database

The image database contains subdirectories for each client that is backed up by NetBackup, including the primary server and any media servers.

The image database is located in the following location:

- **Windows:** `Program Files\Veritas\Netbackup\db\images`
- **UNIX:** `/usr/opensv/netbackup/db/images`

The image database contains the following files:

Image files	Files that store only backup set summary information.
.lck files	Used to prevent simultaneous updates on images.
Image .f files	Used to store the detailed information about each file backup.
db_marker.txt	Used to ensure that access to the db directory is valid when the NetBackup Database Manager starts up. Do not delete this file.

The image database is the largest part of the NetBackup catalog. It consumes about 99% of the total space that is required for the NetBackup catalog. While most of the subdirectories are relatively small in the NetBackup catalogs, `\images` (Windows) or `/images` (UNIX) can grow to hundreds of gigabytes. The image database on the primary server can grow too large to fit on a single tape. Image database growth depends on the number of clients, policy schedules, and the amount of data that is backed up.

See [“Estimating catalog space requirements”](#) on page 668.

If the image catalog becomes too large for the current location, consider moving it to a file system or disk partition that contains more space.

See [“Moving the image catalog”](#) on page 670.

The catalog conversion utility (`cat_convert`) can be used to convert .f files into a human-readable format.

About NetBackup image .f files

The binary catalog contains one or more image .f files. This type of file is also referred to as a “files” file. The image .f file may be large because it contains the detailed backup selection list for each file backup. Generally, image files range in size from 1 kilobyte to 10 gigabytes.

Note: You can use intelligent catalog archiving (ICA) to reduce the number of catalog .f files based on a specified retention period or file size.

See [“Enabling intelligent catalog archiving \(ICA\) to reduce the number of .f files”](#) on page 660.

ICA applies only to servers running NetBackup 11.0 and later using MSDP or MSDP Cloud storage.

The .f files are found in the following location:

Windows: `install_path\NetBackup\db\images\clientname\ctime`

UNIX: `/usr/opensv/netbackup/db/images/clientname/ctime/`

The file layout determines whether the catalog contains one `.f` file or many `.f` files. NetBackup configures the file layout automatically, based on the size of the binary catalog. NetBackup uses one of two layouts: single file layout or multiple file layout.

- Image `.f` file single file layout

NetBackup stores file information in a single image `.f` file if the information for the catalog is less than 100 megabytes.

When the backup file of one catalog backup is less than 100 megabytes, NetBackup stores the information in a single image `.f` file. The image `.f` file is always greater than or equal to 72 bytes, but less than 100 megabytes.

The following is a UNIX example of an `.f` file in a single file layout:

```
-rw----- 1 root other  979483 Aug 29 12:23 test_1030638194_FULL.f
```

- Image `.f` file multiple file layout

When the file information for one catalog backup is greater than 100 megabytes, the information is stored in multiple `.f` files: one main image `.f` file plus nine additional `.f` files.

Separating the additional `.f` files from the image `.f` file and storing the files in the `catstore` directory improves performance while writing to the catalog.

The main image `.f` file is always exactly 72 bytes.

```
-rw- 1 root other      72 Aug 30 00:40 test_1030680524_INCR.f
-rw- 1 root other     804 Aug 30 00:08 catstore/test_1030680524_INCR.f-list
-rw- 1 root other 1489728 Aug 30 00:39 catstore/test_1030680524_INCR.f_imgDir0
-rw- 1 root other      0 Aug 30 00:40 catstore/test_1030680524_INCR.f_imgExtraObj0
-rw- 1 root other 1280176 Aug 30 00:39 catstore/test_1030680524_INCR.f_imgFile0
-rw- 1 root other     192 Aug 30 00:40 catstore/test_1030680524_INCR.f_imgHeader0
-rw- 1 root other      0 Aug 30 00:40 catstore/test_1030680524_INCR.f_imgNDMP0
-rw- 1 root other 9112680 Aug 30 00:39 catstore/test_1030680524_INCR.f_imgRecord0
-rw- 1 root other 2111864 Aug 30 00:39 catstore/test_1030680524_INCR.f_imgStrings0
-rw- 1 root other      11 Aug 30 00:40 catstore/test_1030680524_INCR.f_imgUserGroupNames0
```

About the catalog backup of cloud configuration files

The following cloud configuration files are backed up during the NetBackup catalog backup process:

All `.txt` files in the `meter` directory, which contain intermediate metering data

- `CloudInstance.xml`
- `cloudstore.conf`

- `libstspienencrypt.conf`
- `libstspimetering.conf`
- `libstspithrottling.conf`
- `libstspicloud_provider_name.conf`

All `.conf` files that are specific to the cloud providers that NetBackup supports

The cloud configuration files that are backed up during the catalog backup process reside at the following locations:

Windows	<code>install_path\Veritas\NetBackup\var\global\wmc\cloud</code>
UNIX	<code>/usr/opensv/var/global/wmc/cloud</code>

The files `CloudProvider.xml` and `cacert.pem` are at the following location:

Windows	<code><installed-path>\NetBackup\var\global\cloud</code>
UNIX	<code>/usr/opensv/var/global/cloud/</code>

Note: The `cacert.pem` file is not backed up during the NetBackup catalog backup process.

This `cacert.pem` file is a cloud provider-specific file. This file is installed as part of the NetBackup installation. This file includes the well-known public cloud vendor CA certificates used by NetBackup.

Archiving the catalog and restoring from the catalog archive

Catalog archiving helps administrators solve the kinds of problems that large amounts of catalog data can pose: large catalogs require a greater amount of disk space and can be time-consuming to back up.

Catalog archiving reduces the size of online catalog data by relocating the large catalog `.f` files to secondary storage. NetBackup administration continues to require regularly scheduled catalog backups, but the backups are faster without the large amount of online catalog data.

You can also use intelligent catalog archiving (ICA) to reduce the number of catalog `.f` files from secondary storage. When you enable ICA, any catalog `.f` file that is older than the specified retention period value is removed from the catalog disk.

You can also specify a size value so that any catalog .f file that is greater than or equal to the size value is removed from the catalog disk.

See [“Enabling intelligent catalog archiving \(ICA\) to reduce the number of .f files”](#) on page 660.

Catalog archiving should not be used as a method to reclaim disk space when a catalog file system fills up. In that situation, investigate catalog compression or add disk space to grow the file system.

For additional catalog archiving considerations, see the following topic:

See [“Catalog archiving considerations”](#) on page 667.

To archive the catalog and restore the catalog archive

- 1 Use `bpcatlist` to determine what images are available to be archived.

Running `bpcatlist` alone does not modify any catalog images. Only when the `bpcatlist` output is piped to `bpcatarc` are the `.f` files backed up, and only when the output is piped to `bpcatrm` will the `.f` files be deleted from disk.

To determine what images have `.f` files on disk that can be archived, run the following command. The `catarcid` column indicates whether the `.f` file is not currently backed up (0) or the `catarcid` of the backup of that image.

```
/usr/opensv/netbackup/bin/admincmd/bpcatlist -online
```

To determine what images have been previously archived and removed from disk, run the following command.

```
/usr/opensv/netbackup/bin/admincmd/bpcatlist -offline
```

The catalog commands are described in detail in the following topic:

See [“Catalog archiving commands”](#) on page 665.

Note: If catalog archiving has not been previously run, this command should return: `No entity was found.`

For example, to display all images for a specific client before January 1, 2017, run the following command:

```
bpcatlist -client name -before Jan 1 2017
```

To display the help for the `bpcatlist` command run this command.

```
bpcatlist -help
```

Once the `bpcatlist` output correctly lists all the images that are to be archived or deleted, other commands can be added.

2 Running the catalog archive.

Before running the catalog archive, create a backup policy named **catarc**. The policy is required for the `bpcatarc` command to successfully process images. The name of the policy reflects that the purpose of the schedule is for catalog archiving.

See the following topic for details about configuring the **catarc** policy:

See [“Creating a catalog archiving policy”](#) on page 664.

To run the catalog archive, first run the `bpcatlist` command with the same options used in step 1 to display images. Then pipe the output through `bpcatarc` and `bpcatrm`.

```
bpcatlist -client all -before Jan 1 2017 | bpcatarc | bpcatrm
```

A new job appears in the **Activity Monitor**. The command waits until the backup completes before it returns the prompt. The command reports an error only if the catalog archive fails, otherwise the commands return to the prompt.

The **File List**: section of the Job Details in the **Activity Monitor** displays a list of image files that have been processed. When the job completes with a status 0, the `bpcatrm` command removes the corresponding `.f` files. If the job fails, no catalog `.f` files are removed.

If `bpcatlist` is piped to `bpcatarc` but the results are not piped to `bpcatrm`, the backup occurs but the `.f` files are not removed from disk. The same `bpcatlist` command can then be rerun and piped to `bpcatrm` to remove the `.f` files.

3 Restoring the catalog archive.

To restore the catalog archive, first use the `bpcatlist` command to list the files that need to be restored. Once `bpcatlist` displays the proper files to restore, run the `bpcatres` command to restore the actual files.

To restore all the archived files from step 2, run the following command:

```
bpcatlist -client all -before Jan 1 2017 | bpcatres
```

This command restores all of the catalog archive files before January 1, 2017.

Enabling intelligent catalog archiving (ICA) to reduce the number of `.f` files

Note: Intelligent catalog archiving (ICA) applies only to servers running NetBackup 11.0 and later using MSDP storage.

You can use intelligent catalog archiving (ICA) to reduce the number of catalog .*fil* files based on a specified retention period or file size. When you enable ICA, any catalog .*fil* file that is older than the specified retention period value is removed from the catalog disk. You can also specify a file size value so that any catalog .*fil* file that is greater than or equal to the size value is removed from the catalog disk.

The main advantage of ICA is that it shortens catalog backup time by reducing the number of .*fil* files that need to be backed up if they meet the required criteria:

- The backup image must be older than the configured ICA retention period.
- The .*fil* file must be larger than or equal to the configured ICA minimum size.
- At least one copy of the backup image must be on MSDP storage and have 1 or more true image restore (TIR) fragments.
- Image catalog .*fil* file has not been recalled in the last 24 hours.
- The backup image must be from a completed SLP or from a backup that is not managed by an SLP.
- The backup image is not from a catalog backup.
- The image catalog is not archived.

When ICA is enabled, you may notice the following behaviors:

- Initial image cleanup after you enable ICA may take longer than usual.
- Catalog backups will be faster if any of the .*fil* files that are involved have been intelligently archived.
- Browse and Restore functions will take longer if any of the .*fil* files that are involved have been intelligently archived.

No additional action is needed to restore the catalog .*fil* file. Catalog .*fil* files are restored from images automatically as follows:

- When an ICA image is browsed.
- When an ICA-eligible copy is expired from an ICA image. Restoring catalog .*fil* files ensures that the remaining copies from that image are accessible and usable.
- When an ICA-eligible image is found but its catalog .*fil* file missing.

More information about .*fil* files is available:

See [“About NetBackup image .*fil* files”](#) on page 655.

To enable intelligent catalog archiving (ICA) and specify retention and file size values

- 1 Run the following command on the primary server:

```
bpconfig -ica_retention seconds
```

When the *seconds* value is between 1 and 2147472000, ICA is enabled. Any image which is older than the value is processed for ICA. The catalog .f file from the ICA-eligible image is removed from the catalog disk. Setting this value to 0 (zero) disables ICA. The default value for NetBackup Flex Scale and Cloud Scale environments is 2592000 (30 days). The default value for all other NetBackup environments is 0 (disabled).

For Accelerator-enabled backups, specify an ICA retention value that is longer than full backup schedules so that the number of .f file restores from ICA images goes down.

For example, to set the ICA retention value to 30 days, enter `bpconfig -ica_retention 2592000`.

Use `bpconfig -U` to verify the change:

```
# bpconfig -U
Admin Mail Address:          sasquatch@wapati.edu
Job Retry Delay:             10 minutes
Max Simultaneous Jobs/Client: 1
Backup Tries:                1 time(s) in 12 hour(s)
Keep Error/Debug Logs:      3 days
Max drives this master:      0
Keep TrueImageRecovery Info: 24 days
Compress DB Files:           (not enabled)
Media Mount Timeout:         30 minutes
Display Reports:             24 hours ago
Preprocess Interval:         0 hours
Image DB Cleanup Interval:    12 hours
Image DB Cleanup Wait Time:   10 minutes
Policy Update Interval:      10 minutes
Intelligent Catalog Archiving: Files file larger than 1024 KB
Intelligent Catalog Archiving: Images older than 30 day(s)
```

-
- 2** **Note:** After you enable ICA, the minimum file size for `.f` files is set to the default value 1024 KB. Use this step to change that value.
-

To specify a minimum file size, run the following command on the primary server:

```
bpconfig -ica_min_size size
```

When the `size` value is between 0 and 2097151, any catalog `.f` file that is larger than or equal to the `size` value is removed from the catalog disk. The default value is 1024.

For example to set the ICA minimum file size to 2048 KB, enter `bpconfig -ica_min_size 2048`.

Use `bpconfig -U` to verify the change:

```
# bpconfig -U
Admin Mail Address:          sasquatch@wapati.edu
Job Retry Delay:             10 minutes
Max Simultaneous Jobs/Client: 1
Backup Tries:                1 time(s) in 12 hour(s)
Keep Error/Debug Logs:      3 days
Max drives this master:      0
Keep TrueImageRecovery Info: 24 days
Compress DB Files:           (not enabled)
Media Mount Timeout:         30 minutes
Display Reports:             24 hours ago
Preprocess Interval:         0 hours
Image DB Cleanup Interval:    12 hours
Image DB Cleanup Wait Time:   10 minutes
Policy Update Interval:       10 minutes
Intelligent Catalog Archiving: Files file larger than 2048 KB
Intelligent Catalog Archiving: Images older than 30 day(s)
```

To disable intelligent catalog archiving (ICA)

- ◆ Run the following command on the primary server:

```
bpconfig -ica_retention 0
```

Use `bpconfig -U` to verify the change:

```
# bpconfig -U
Admin Mail Address:          sasquatch@wapati.edu
Job Retry Delay:             10 minutes
Max Simultaneous Jobs/Client: 1
Backup Tries:                1 time(s) in 12 hour(s)
Keep Error/Debug Logs:      3 days
Max drives this master:      0
Keep TrueImageRecovery Info: 24 days
Compress DB Files:           (not enabled)
Media Mount Timeout:         30 minutes
Display Reports:             24 hours ago
Preprocess Interval:         0 hours
Image DB Cleanup Interval:    12 hours
Image DB Cleanup Wait Time:   10 minutes
Policy Update Interval:       10 minutes
Intelligent Catalog Archiving: (not enabled)
```

Creating a catalog archiving policy

The catalog archiving feature requires the presence of a policy named **catarc** before the catalog archiving commands can run properly. The policy can be reused for catalog archiving.

To create a catalog archiving policy

- 1 Open the NetBackup web UI.
- 2 On the left, click **Protection > Policies**. Then click **Add**.
- 3 Enter the **Policy name catarc**.

The **catarc** policy waits until `bpcatarc` can activate it. Users do not run this policy. Instead, `bpcatarc` activates this special policy to perform a catalog backup job, then deactivates the policy after the job is done.

- 4 In the **Attributes** policy tab, set the **Policy type** to **Standard** or **MS-Windows**, according to the platform of the primary server.
- 5 In the **Attributes** policy tab, deactivate the catalog archive policy by clearing the **Go into effect at** box.

- 6 Select the **Schedules** tab and click **Add** to create a schedule.
 In the **Attributes** schedule tab, the **Name** of the schedule is not restricted, but the **Type of backup** must be **User backup**.
- 7 Select a **Retention** for the catalog archive. Set the retention level for a time at least as long as the longest retention period of the backups being archived.
 Data can be lost if the retention level of the catalog archive is not long enough.
 You may find it useful to set up and then designate a special retention level for catalog archive images.
- 8 Select the **Start window** tab and define a schedule for the **catarc** policy.
 The schedule must include in its window the time when the `bpcatarc` command is run. If the `bpcatarc` command is run outside of the schedule, the operation fails.
- 9 Click **Add** to save the schedule.
- 10 On the **Clients** tab, enter the name of the primary server as it appears on the NetBackup servers list.
- 11 On the **Backup selections** tab, browse to the directory where catalog backup images are placed:
 On Windows: `install_path\NetBackup\db\images`
 On UNIX: `/usr/openv/netbackup/db/images`
- 12 Click **Create** to save the policy.

Catalog archiving commands

The catalog archiving option relies on three commands to designate a list of catalog `.f` files, then archive the files. A fourth command, `bpcatres`, is used to restore the files if necessary.

Catalog archiving uses the following commands.

Table 50-1 Catalog archiving commands

Command	Description
bpcatlist	The <code>bpcatlist</code> command queries the catalog data. Then, <code>bpcatlist</code> lists the portions of the catalog that are based on selected parameters. For example, date, client, policy, schedule name, backup ID, the age of the backup image, or the date range of the backup image. <code>bpcatlist</code> outputs the formatted image summary information of matched images to standard output. The other catalog archiving commands, <code>bpcatarc</code>, <code>bpcatrm</code>, and <code>bpcatres</code>, all depend on input from <code>bpcatlist</code> by a piped command. For example, to archive (backup and delete) all of the <code>.f</code> files that were created before January 1, 2012, the following would be entered: <pre>bpcatlist -client all -before Jan 1 2012 bpcatarc bpcatrm</pre> <code>bpcatlist</code> is also used to provide status information. For each catalog, it lists the following information: ■ Backup ID (Backupid) ■ Backup date (Backup Date) ■ Catalog archive ID (catarcid). After one <code>.f</code> file is successfully backed up, a catalog archive ID is entered into the catarcid field in the image file. This field is zero (0) if the image was never archived. ■ Archived status (S). Indicates whether the catalog was archived (2) or was not archived (1). ■ Compressed status (C). Indicates whether the catalog was compressed (<i>positive_value</i>) or was not compressed (0). ■ Catalog file name (Files file) The following is an example of the <code>bpcatlist</code> output, showing all of the backups for client alpha since October 23: <pre># bpcatlist -client alpha -since Oct 23 Backupid Backup Date ...Catarcid S C Files file alpha_097238 Oct 24 10:47:12 2012 ... 973187218 1 0 alpha_097238_UBAK.f alpha_097233 Oct 23 22:32:56 2012 ... 973187218 1 0 alpha_097233_FULL.f alpha_097232 Oct 23 19:53:17 2012 ... 973187218 1 0 alpha_097232_UBAK.f</pre> More information is available in the NetBackup Commands Reference Guide.
bpcatarc	The <code>bpcatarc</code> command reads the output from <code>bpcatlist</code> and backs up the selected list of <code>.f</code> files. After one <code>.f</code> file is successfully backed up, a catalog archive ID is entered into the catarcid field in the image file. For archiving of the <code>.f</code> files to proceed, a policy by the name of catarc is required. The policy is based on a User Backup type schedule. The schedule for catarc must include in its window the time <code>bpcatarc</code> command is run. See “Creating a catalog archiving policy” on page 664.

Table 50-1 Catalog archiving commands (*continued*)

Command	Description
<code>bpcatrm</code>	The <code>bpcatrm</code> command reads the output from <code>bpcatlist</code> or <code>bpcatarc</code>. If the image file has valid catarcid entries, <code>bpcatrm</code> deletes selected image .f files from the online catalog. <code>bpcatrm</code> does not remove one .f file unless the file has been previously backed up using the catarc policy.
<code>bpcatres</code>	Use the <code>bpcatres</code> command to restore the catalog. The <code>bpcatres</code> command reads the output from <code>bpcatlist</code> and restores selected archived .f files to the catalog. For example: <pre>bpcatlist -client all -before Jan 1 2012 bpcatres</pre>

Catalog archiving considerations

Consider the following items before catalog archiving:

- Perform catalog archiving operations when NetBackup is in an inactive state (no jobs are running).
- Catalog archiving modifies existing catalog images. As a result, it should never be run when the catalog file system is 100% full.
- To ensure that catalog backup images are not on the same tapes as user backups, create a separate media pool for catalog archives.
- You may find it useful to set up and then designate, a special retention level for catalog archive images.
To specify retention levels, open the NetBackup web UI. On the left click **Hosts > Host properties**. Locate the primary server and click **Edit primary server**. Then click **Retention periods**.
- Additional time is required to mount the tape and perform the restore of archived .f files.
- There is no simple method to determine to which tape the catalog has been archived. The `bpcatlist -offline` command is the only administrative command to determine what images have been archived. This command does not list what tape was used for the archive. As a result, exercise caution to ensure that the tapes used for catalog archiving are available for restoring the archived catalog images. Either create a separate volume pool to use exclusively for catalog archives or find a method to label the tape as a catalog archive tape.

Extracting images from the catalog archives

The situation may arise in which a storage provider needs to extract all of a specific client's records. The storage provider can extract the customer images from the catalog archive by creating the archives that are based on client name.

To extract images from the catalog archives based on a specific client

- 1 Create a volume pool for the client.
- 2 Create a catalog archiving policy. Indicate the volume pool for that client in the **Attributes** tab.
- 3 Run `bpcatlist` so only the `.f` files from that client are listed. For example:

```
bpcatlist -client clientname | bpcatarc | bpcatrm
```
- 4 If you do not want to write more images to the client's volume pool, change the volume pool before you run another archiving catalog.

Estimating catalog space requirements

NetBackup requires disk space to store its error logs and information about the files it backs up.

The disk space that NetBackup needs varies according to the following factors:

- Number of files to be backed up
- Frequency of full and incremental backups
- Number of user backups and archives
- Retention period of backups
- Average length of full path of files
- File information (such as owner permissions)
- Average amount of error log information existing at any given time
- Whether you have enabled the database compression option.

To estimate the disk space that is required for a catalog backup

- 1 Estimate the maximum number of files that each schedule for each policy backs up during a single backup of all its clients.
- 2 Determine the frequency and the retention period of the full and the incremental backups for each policy.

- 3 Use the information from steps 1 and 2 to calculate the maximum number of files that exist at any given time.

For example:

Assume that you schedule full backups to occur every seven days. The full backups have a retention period of four weeks. Differential incremental backups are scheduled to run daily and have a retention period of one week.

The number of file paths you must allow space for is four times the number of files in a full backup. Add to that number one week's worth of incremental backups.

The following formula expresses the maximum number of files that can exist for each type of backup (daily or weekly, for example):

Files per Backup × Backups per Retention Period = Max Files

For example:

A daily differential incremental schedule backs up 1200 files and the retention period for the backup is seven days. Given this information, the maximum number of files that can exist at one time are the following:

$$1200 \times 7 \text{ days} = 8400$$

A weekly full backup schedule backs up 3000 files. The retention period is four weeks. The maximum number of files that can exist at one time are the following:

$$3000 \times 4 \text{ weeks} = 12,000$$

Obtain the total for a server by adding the maximum files for all the schedules together. Add the separate totals to get the maximum number of files that can exist at one time. For example, 20,400.

For the policies that collect true image restore information, an incremental backup collects catalog information on all files (as if it were a full backup). This changes the calculation in the example: the incremental changes from $1200 \times 7 = 8400$ to $3000 \times 7 = 21,000$. After 12,000 is added for the full backups, the total for the two schedules is 33,000 rather than 20,400.

- 4 Obtain the number of bytes by multiplying the number of files by the average number of bytes per file record.

If you are unsure of the average number of bytes per file record, use 132. The results from the examples in step 3 yield:

$$(8400 \times 132) + (12,000 \times 132) = 2692800 \text{ bytes (or about 2630 kilobytes)}$$

- 5 Add between 10 megabytes to 15 megabytes to the total sum that was calculated in step 4. The additional megabytes account for the average space that is required for the error logs. Increase the value if you anticipate problems.
- 6 Allocate space so all the data remains in a single partition.

NetBackup file size considerations on UNIX systems

File system limitations on UNIX include the following:

- Some UNIX systems have a large file support flag. Turn on the flag to enable large file support.
- Set the file size limit for the root user account to unlimited to support large file support.

Moving the image catalog

An image catalog may become too large for its current location. Consider moving the image catalog to a file system or disk partition that contains more available space.

Notes about moving the image catalog

- NetBackup does not support saving the catalog to a remote NFS share. CIFS is supported on some SAN or NAS storage.
See [“Parts of the NetBackup catalog”](#) on page 651.
- NetBackup only supports moving the image catalog to a different file system or disk partition. It does not support moving the other subdirectories that make up the entire NetBackup catalog.
For example, on Windows, do not use the `ALTPATH` mechanism to move `install_path\NetBackup\db\error`.
For example, on UNIX, do not move `/usr/opensv/netbackup/db/error`. The catalog backup only follows the symbolic link when backing up the `/images` directory. So, if symbolic links are used for other parts of the NetBackup catalog, the files in those parts are not included in the catalog backup.
- The directory that is specified in the `ALTPATH` file is not automatically removed if NetBackup is uninstalled. If NetBackup is uninstalled, you must manually remove the contents of this directory.

Moving the image catalog between Windows hosts

To move the image catalog on Windows

- 1 Back up the NetBackup catalogs manually.

A backup of the catalogs ensures that you can recover image information in case something is accidentally lost during the move.

See [“Backing up NetBackup catalogs manually”](#) on page 386.

- 2 Check the **Jobs** tab in the **Activity monitor** and ensure that no backups or restores are running for the client.

If jobs are running, either wait for them to end or stop them by using the **Jobs** tab in the Activity monitor.

- 3 Use the **Daemons** tab in the **Activity monitor** to stop the Request Manager and the Database Manager daemons. These services are stopped to prevent jobs from starting. Do not modify the database while this procedure is performed.

- 4 Create a file named `ALTPATH` in the image catalog directory.

For example, if NetBackup is installed in the default location and the client name is *mars*, the path to the image catalog is:

```
C:\Program Files\Veritas\NetBackup\db\images\mars\ALTPATH
```

- 5 Create the directory to which you intend to move the image information. For example:

```
E:\NetBackup\alternate_db\images\client_name
```

- 6 On the first line of the `ALTPATH` file, specify the path to the directory where you intend to move the client's image information. For example:

```
E:\NetBackup\alternate_db\images\client_name
```

The path is the only entry in the `ALTPATH` file.

- 7 Move all files and directories (except the `ALTPATH` file) that are in the current client directory to the new directory.

For example, if the images are currently in

```
C:\Program Files\Veritas\NetBackup\db\images\mars
```

and the `ALTPATH` file specifies

```
E:\NetBackup\alternate_db\images\mars
```

then move all files and directories (except the `ALTPATH` file) to

```
E:\NetBackup\alternate_db\images\mars
```

- 8 Start the NetBackup Request Daemon, NetBackup Job Manager, and NetBackup Policy Execution manager in the **Daemons** tab.

Backups and restores can now resume for the client.

Moving the image catalog between UNIX hosts

To move the image catalog on UNIX

- 1 Check that no backups are in progress by running:

```
/usr/opensv/netbackup/bin/bpps
```

- 2 Stop `bprd` by running:

```
/usr/opensv/netbackup/bin/admincmd/bprdreq -terminate
```

- 3 Stop `bpdbm` by running:

```
/usr/opensv/netbackup/bin/bpdbm -terminate
```

- 4 Create the directory in the new file system. For example:

```
mkdir /disk3/netbackup/db/images
```

- 5 Move the image catalog to the new location in the other file system.

- 6 Create a symbolic link from `/usr/opensv/netbackup/db/images` to the new location in the other file system.

See [“NetBackup file size considerations on UNIX systems”](#) on page 670.

About image catalog compression

The image catalog contains information about all client backups. It is accessed any time a user lists or restores files. NetBackup lets you compress all portions of the catalog or only older portions of the catalog.

Control image catalog compression by setting the **Compress catalog interval** in the **Global attributes** host property. This interval indicates how old the backup

information must be before it is compressed. Specify the number of days to defer compression information, so users who restore files from recent backups are not affected. By default, **Compress catalog interval** is set to 0 and image compression is not enabled.

Note: Cohesity discourages manually compressing or decompressing the catalog backups with the `bpimage -[de]compress` command or any other method. Manually compressing or decompressing a catalog backup while any backup (regular or catalog) is running results in inconsistent image catalog entries. When users list and restore files, the results can be incorrect.

It does not make a difference to NetBackup if the backup session was successful. The operation occurs while NetBackup expires backups and before it runs the `session_notify` script and the backup of the NetBackup catalogs.

The time to perform compression depends on the server speed and the number and size of the files being compressed. Files are compressed serially, and temporary working space is required in the same partition.

When numerous compressed image catalog files must be processed, the backup session is extended until compression is complete. The additional backup time is especially noticeable the first time you perform the compression. To minimize the effect of the initial sessions, consider compressing the files in stages. For example, begin by compressing the records for the backups older than 120 days. Continue to reduce the number of days over a period of time until you reach a comfortable setting.

Compressing the image catalog accomplishes the following objectives:

- Reduces greatly the disk space that is consumed.
- Reduces the media that is required to back up the catalog.

The amount of space that is reclaimed varies with the types of backups you perform. Full backups result in a larger percentage of catalog compression than incremental backups. Normally, more data is duplicated in a catalog file for a full backup. Using catalog compression, a reduction of 80% is possible.

This reduction in disk space and media requirements is achieved at the expense of performance when a user lists or restores files. Since the information is uncompressed at each reference, performance degradation is in direct proportion to the number and size of compressed files that are referenced. If the restore requires numerous catalog files to be uncompressed, increase the **File browse timeout** value that is associated with list requests. (See the **Timeouts** host property for the client.)

Uncompressing the NetBackup catalog

You may find it necessary to temporarily uncompress all records that are associated with an individual client. Uncompress the records if you anticipate large or numerous restore requests, for example.

To uncompress the NetBackup catalog on Windows

- 1 Verify that the partition where the image catalog resides contains enough space to accommodate the uncompressed catalog.
See [“Estimating catalog space requirements”](#) on page 668.
- 2 Stop the NetBackup Request Daemon service, `bprd`.
- 3 Verify that the NetBackup Database Manager, `bpdbm`, is running.
- 4 In the NetBackup web UI, select **Hosts > Host properties**.
- 5 Select the primary server and click **Connect**. Then select the server and click **Edit primary server**.

- 6 Select **Global attributes**.

- 7 Clear the **Compress catalog interval** check box. Then click **Save**.

- 8 Open a command prompt. Change to the following directory:

```
install_path\Veritas\NetBackup\bin\admincmd
```

Run one of the followings commands.

To decompress the records for a specific client, enter:

```
bpimage -decompress -client_name
```

To decompress the records for all clients, enter:

```
bpimage -decompress -allclients
```

- 9 Restart the NetBackup Request Daemon (`bprd`).
- 10 Restore the files from the client.
- 11 Set the **Compress catalog interval** to its previous value.

The records that were uncompressed for this client are compressed after the next backup schedule.

To uncompress the NetBackup catalog on UNIX

- 1 Perform the following steps as root on the primary server to uncompress the NetBackup catalog.

Verify that the partition where the image catalog resides has enough space to uncompress the client's image records.

- 2 Stop the request daemon, `bprd`, by running:

```
/usr/opensv/netbackup/bin/admincmd/bprdreq -terminate
```

- 3 Make sure that `bpdbm` is running:

```
/usr/opensv/netbackup/bin/bpps
```

- 4 In the NetBackup web UI, select **Hosts > Host properties**.

- 5 Select the primary server and click **Connect**. Then select the server and click **Edit primary server**.

- 6 Select **Global attributes**.

- 7 Clear the **Compress catalog interval** check box. Then click **Save**.

- 8 Change your working directory to `/usr/opensv/netbackup/bin` and run the command:

```
admincmd/bpimage -decompress -client name
```

- 9 Restart the request daemon `bprd`. Run the following command:

```
/usr/opensv/netbackup/bin/initbprd
```

- 10 Restore the files from the client.

- 11 Set the **Compress catalog interval** to its previous value.

The records that were uncompresssed for this client are compressed after the next backup schedule.

About the file hash search in NetBackup

You can search the files using the file hash of the file. The file hash search is supported on the accelerator backups. After you enable this feature, the SHA-256 hash of the files is calculated during the client backups and saved to the NetBackup primary server catalog.

The file hash calculation is supported only for normal files. The sparse file, symbolic link, and any other special files are not supported.

Perform the following steps to configure and use the feature:

1. Configure file hash server.

See “Configuring the file hash server” on page 676.

2. Enable Calculating the file hash.

See “Calculating the file hash” on page 677.

3. Search for files using the file hash.

See “Searching the files using the file hash” on page 678.

The NetBackup primary server, media server, and a client must be 10.5 or later to use this feature.

Configuring the file hash server

The file hash server is used to store and search all file hashes in one NetBackup domain. You must configure the file hash server first and then configure the backup policy to calculate the file hash.

See “Calculating the file hash” on page 677.

Ensure that the following prerequisites are met if the file hash server is not installed with MSDP storage server:

- Ensure that NGINX is installed and running on the server. The minimum recommended NGINX version is 1.24.0.
- If SE Linux is configured, ensure that the `polycoreutils` and `polycoreutils-python` (for RHEL 7) or `polycoreutils-python-utils` (for RHEL 8) packages are installed from the same RHEL Yum source (RHEL server) and then run the following commands:

```
semanage port -a -t http_port_t -p tcp 10087
setsebool -P httpd_can_network_connect 1
```

- Enable the logrotate permission in SELinux using the following command:

```
semanage permissive -a logrotate_t
```

The file hash server supports only the security certificates that the NetBackup certificate authority issues.

The file hash server name must be the same as the NetBackup media server, which is the NetBackup host name in the `bp.conf` file.

To configure the file hash server

- ◆ Run the following command to configure file hash server on the media server.

```
/usr/opensv/pdde/pdcr/bin/fhdb_config.sh
--hash-storage-path=<hash_db_path>
```

Enabling the file hash server on the NetBackup primary server

After you configure the file hash server, you must enable it on the NetBackup primary server.

To enable the file hash server on the NetBackup primary server

- ◆ Run the following command to configure file hash server on the primary server.

```
/usr/opensv/netbackup/bin/goodies/nbfhsmgr -config <file hash  
hostname>
```

Calculating the file hash

In the backup policy, enable the option **Use Accelerator** and then enable the option **Calculate file hash** to calculate the SHA-256 of the files. After SHA-256 is calculated, the file hash information is saved to the NetBackup catalog. The file hash information is then copied to the file hash server if the file hash server is configured.

The NetBackup catalog is expected to increase by 20% or more. You can configure to delete the file hash information automatically from the NetBackup catalog after it is copied to the file hash server. Add the line

`AUTO_CLEAN_FILE_HASH_FROM_CATALOG = 1` in the `bp.conf` file on the primary server.

You can calculate the file hash for the following policy types:

- Windows
- Standard
- NAS-Data-Protection

Note: This feature may affect the backup performance depending on the client configuration such as CPU and memory. If the CPU has SHA extensions, the hash calculation is faster than the CPU without SHA extensions.

To enable calculating the file hash

- 1 On the left, click **Protection > Policies** and then click **Add** to add a new policy or select the existing policy to edit it.
- 2 On the **Attributes** tab, select **Use Accelerator** and then select **Calculate file hash**.

After you enable this feature, the next backup is a full backup. It treats all data as changed data because it calculates the file hash for each file. The subsequent backups calculate the file hash for changed files only.

- 3 Click **Create** or **Save**.

Searching the files using the file hash

You can search the files using the file hash. You must configure the file hash server and enable the **Calculate file hash** option in the policy on the Web UI. After this option is enabled, SHA-256 information is saved to the NetBackup catalog and the file hash server.

To search the files using the file hash

- 1 On the left, click **Catalog**.
- 2 On the **Search** tab, select **File hash search** from the **Actions** list.
- 3 You can limit the number of file hash searches by navigating to **Manage file hash search settings** on the top right corner and providing the limit value in **Limit number of file hash searches**. Default value is 50.
- 4 You can add a tag for your search in **Tag this search**.

Later, you can filter the file hash search results with this tag.

- 5 Enter the list of SHA-256 hash strings of the files to search the files.

It searches the entire image catalog, which may take some time. You can see the job status and results in the Activity monitor.

Setting the resource limit for file hash search

User can set the maximum number of jobs that can run on file hash server.

If a new search request is received after exceeding the resource limit that was set, then file hash search job will go to queued state. The job would be activated only after the earlier jobs are completed.

Identifying the backups that have the file hash enabled

You can list the backups that have the file hash enabled. The file hash is calculated for these backups and saved to the NetBackup catalog.

To identify the backups that have the file hash enabled.

- ◆ To list all the backups that have file hash enabled, run the following command:

```
/usr/opensv/netbackup/bin/admincmd/bpcatlist -file-hash-present
```

Removing the file hash from the backup

You can remove the file hash from the backup if you do not want the file hash data for the specific backup. The file hash data is removed from the NetBackup catalog.

To remove the file hash from the backup

- ◆ To remove the file hash from the backup, run the following command:

```
/usr/opensv/netbackup/bin/admincmd/bpimage -removehash -backupid  
bid
```

Migrating the file hash data from one server to another

You might need to migrate the file hash data from one server to another when disk is full or for any other reasons.

To migrate the hash from one server to another

- 1 Stop the `crond` and `spws` services on the existing server:

```
service crond stop /usr/opensv/pdde/vpfs/etc/init.d/pdde-spws stop
```

- 2 Configure the file hash service on the new server.

See [“Configuring the file hash server”](#) on page 676.

- 3 Stop the `crond` and `spws` services on the new server:

```
service crond stop /usr/opensv/pdde/vpfs/etc/init.d/pdde-spws stop
```

- 4 Copy the following folders from the existing file hash server to

<new_file_hash_storage_path>/hashdb/ on the new file hash server:

```
<existing_file_hash_storage_path>/hashdb/recv  
< existing_file_hash_storage_path>/hashdb/catalog  
< existing_file_hash_storage_path>/hashdb/sqlite  
< existing_file_hash_storage_path>/hashdb/history
```

- 5 Change the folder's owner to file hash service user on the new file hash server.

- Get the service user on the new file hash server.

```
cat /usr/opensv/netbackup/bp.conf |grep "^SERVICE_USER"
```
 - Change the owner.

```
chown -Rv ${service_user} <new_file_hash_storage_path>/hashdb/
```
- 6** Start the `cron` and `spws` services on the new server:
- ```
service cron start /usr/opensv/pdde/vpfs/etc/init.d/pdde-spws
start
```
- 7** Enable the new file hash server on the NetBackup primary server.  
See [“Enabling the file hash server on the NetBackup primary server”](#) on page 677.

---

**Note:** After you enable the new file hash server on the primary server, the existing file hash server will no longer be used.

---

## Configuring the backup of file hash data on the file hash server

We recommend that you take a backup of the file hash data regularly. You can use the Standard NetBackup policy type to back up the file hash data, which you can recover in the event of a disaster.

See [“Restoring the file hash data to the file hash server”](#) on page 681.

### To back up the file hash data

- 1** On the left, select **Protection > Policies**.
- 2** Click **Add**.
- 3** On the **Attributes** tab, select **Standard** policy type.
- 4** On the **Attributes** tab, select **Collect true image restore information** to enable true image restore (TIR).
- 5** On the **Schedules** tab, configure all the necessary schedules.

We recommend that you schedule the daily backup because the file hash data is updated on the server once a day.

- 6** On the **Clients** tab, add the file hash server as a client.



- 7 On the **Backup selections** tab, add the following directories in the order:

```
<file_hash_storage_path>/hashdb/history
<file_hash_storage_path>/hashdb/recv
<file_hash_storage_path>/hashdb/sqlite
<file_hash_storage_path>/hashdb/catalog
```

- 8 Click **Create**.

## Restoring the file hash data to the file hash server

You can restore the file hash data to the file hash server in the event of a disaster.

### To restore the file hash data to the file hash server

- 1 If the file hash service is not running on the file hash server, you must reconfigure the file hash server.

See [“Configuring the file hash server”](#) on page 676.

- 2 If the file hash service is running on the file hash server, stop the `crond` and `spws` services on the server:

```
service crond stop /usr/opensv/pdde/vpfs/etc/init.d/pdde-spws stop
```

- 3 Restore the following directories:

```
<file_hash_storage_path>/hashdb/history
<file_hash_storage_path>/hashdb/recv
<file_hash_storage_path>/hashdb/sqlite
<file_hash_storage_path>/hashdb/catalog
```

- 4 Start the `crond` and `spws` services on the server:

```
service crond stop /usr/opensv/pdde/vpfs/etc/init.d/pdde-spws stop
```

# About the NetBackup database

This chapter includes the following topics:

- [About the NetBackup database installation](#)
- [Post-installation tasks](#)
- [Using the NetBackup Database Administration utility on Windows](#)
- [Using the NetBackup Database Administration utility on UNIX](#)

## About the NetBackup database installation

Generally, the implementation of the NetBackup database in the NetBackup catalog is transparent. The NetBackup primary server includes a private, non-shared database server for the NetBackup database (NBDB).

The same installation of the NetBackup database is used for the optionally-licensed product, Bare Metal Restore (BMR) and its associated database (BMRDB). The BMR database is created during the BMR installation process.

By default, the NetBackup database (NBDB) is installed on the primary server. The primary server is also the default location for the Enterprise Media Manager (EMM). Since EMM is the primary user of NBDB, the NetBackup database always resides on the same computer as the Enterprise Media Manager.

See [“About the Enterprise Media Manager \(EMM\)”](#) on page 653.

## About NetBackup primary server installed directories and files

The NetBackup Scale-Out Relational Database is installed in the following directories.

## Windows

`install_path\Veritas\NetBackupDB`

`install_path\Veritas\NetBackup\bin`

`install_path\Veritas\NetBackupDB\data\instance`

The databases are installed in the following subdirectories:

`install_path\Veritas\NetBackupDB\data\nbdb\`

`install_path\Veritas\NetBackupDB\data\nbazdb\`

`install_path\Veritas\NetBackupDB\data\bmrdb\` (if BMR is installed)

## On UNIX

`/usr/opensv/db`

`/usr/opensv/var/global`

`/usr/opensv/db/data/instance/`

The databases are installed in the following subdirectories:

`/usr/opensv/db/data/nbdb/`

`/usr/opensv/db/data/nbazdb/`

`/usr/opensv/db/data/bmrdb/`

## About the `bin` directory

The `bin` is located as follows:

`install_path\Veritas\NetBackup\bin`

---

**Warning:** Use these utilities and commands in this directory with caution.

---

Contains the utilities and binaries for running and administering NetBackup services. More information can be found in the *NetBackup Commands Reference Guide*.

For information on using the NetBackup Database Administration utility (`NbDbAdmin.exe` or `dbadm`), see the following topics:

See [“Using the NetBackup Database Administration utility on Windows”](#) on page 693.

See [“Using the NetBackup Database Administration utility on UNIX”](#) on page 698.

## About the contents of the NetBackupDB and db directories

The following table describes the contents of the following directories.

On Windows: *install\_path*\Veritas\NetBackupDB\

On UNIX: */usr/openv/db/*

**Table 51-1** NetBackupDB and db directory contents

| Directory | Description                                                                                                                                                                                                                             |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| bin       | Contains the utilities and commands for administering the NetBackup database service.                                                                                                                                                   |
| data      | The default location of the NetBackup databases (NBDB, NBAZDB, and BMRDB) and certain configuration files.                                                                                                                              |
| lib       | On UNIX: Contains all the shared libraries for the NetBackup Scale-Out Relational Database. The directory also includes ODBC libraries, used to connect to NBDB and BMRDB.                                                              |
| scripts   | <b>Warning:</b> Do not edit the scripts that are located in this directory.<br><br>Contains the scripts that are used to create the NetBackup database. It also contains the scripts that are used to create the EMM and other schemas. |
| share     | Contains the PostgreSQL document and module files that are required by the NetBackup database server.                                                                                                                                   |
| staging   | Used as a temporary staging area during catalog backup and recovery.                                                                                                                                                                    |
| WIN64     | (Windows) Contains .dll files for the NetBackup Scale-Out Relational Database.                                                                                                                                                          |

## About the data directory

The following directory is the default location of the NetBackup database, NBDB:

On Windows: *install\_path*\NetBackupDB\data

On UNIX: */usr/openv/db/data*

The *\data\* directory contains the following subdirectories and files:

- *bmrdb*  
If BMR is installed, this directory contains the BMR database.
- *nbdb*  
The main NetBackup database, including EMM.
- *nbazdb*  
The NetBackup Authorization database.
- *vxdbs.conf*

The file that contains the configuration information specific to the installation of the NetBackup database.

See “[vxdbms.conf](#)” on page 685.

- `nbdbinfo.dat`  
A backup of the NetBackup DBA password.

## vxdbms.conf

On Windows:

```
VXDBMS_NB_SERVER = NB_server_name
VXDBMS_NB_DATABASE = NBDB
VXDBMS_BMR_DATABASE = BMRDB
VXDBMS_AZ_DATABASE = NBAZDB
VXDBMS_NB_STAGING = C:\Program Files\Veritas\NetBackupDB\staging
VXDBMS_NB_PORT = 13785
VXDBMS_NB_DATA = C:\Program Files\Veritas\NetBackupDB\data
VXDBMS_NB_PASSWORD = encrypted_password
AZ_DB_PASSWORD = encrypted_password
VXDBMS_POSTGRESQL_POOLER_ODBC_PORT = 13787
```

On UNIX:

```
VXDBMS_NB_SERVER = NB_server_name
VXDBMS_NB_PORT = 13785
VXDBMS_NB_DATABASE = NBDB
VXDBMS_AZ_DATABASE = NBAZDB
VXDBMS_BMR_DATABASE = BMRDB
VXDBMS_NB_DATA = /usr/opensv/db/data
VXDBMS_NB_STAGING = /usr/opensv/db/staging
VXDBMS_NB_PASSWORD = encrypted_password
AZ_DB_PASSWORD = encrypted_password
VXDBMS_POSTGRESQL_POOLER_ODBC_PORT = 13787
```

The encrypted password that is used to log into the DBA accounts is stored in `vxdbms.conf`. These accounts include NBDB, NBAZDB, and BMRDB and other data accounts.

## NetBackup configuration entry

The `VXDBMS_NB_DATA` registry entry (Windows) or the `bp.conf` entry (UNIX) is a required entry and is created upon installation. The entry indicates the path to the directory where the following are located: NetBackup database, authorization database, BMR database, and the `vxdbms.conf` file.

#### On Windows:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Veritas\NetBackup\CurrentVersion\
Config\VXDBMS_NB_DATA
```

On UNIX: /usr/openv/netbackup/bp.conf

```
VXDBMS_NB_DATA = /usr/openv/db/data
```

## NetBackup database server management

This topic describes the commands that are available to manage the NetBackup database.

To start and stop the NetBackup database, use one of the following methods:

- In the **Daemons** tab of the Activity monitor, select the service **NetBackup Scale-Out Relational Database Manager** (vrtsdbsvc\_psql).
- (Windows) From the Windows Service Manager, select the service **NetBackup Scale-Out Relational Database Manager** (vrtsdbsvc\_psql).

- (Windows) Use the following commands:

```
install_path\Veritas\NetBackup\bin\bpdown -e vrtsdbsvc_psql
```

- `install_path\Veritas\NetBackup\bin\bpup -e vrtsdbsvc_psql`

- (UNIX) Use the following commands:

```
/usr/openv/db/bin/nbdbms_start_server -start
```

Starts the NetBackup Scale-Out Relational Database server if no option is specified.

```
/usr/openv/db/bin/nbdbms_start_server -stop -f
```

Stops the server; `-f` forces a shutdown with active connections.

The **NetBackup Scale-Out Relational Database Manager** daemon is included in the `stop` command or the `start` command, which starts and stops all NetBackup daemons.

Individual databases can be started or stopped, while the NetBackup Scale-Out Relational Database Manager service continues. Use the NetBackup Database Administration utility or the following commands:

- `nbdb_admin [-start | -stop]`

Starts or stops NBDB without shutting down the NetBackup Scale-Out Relational Database server.

To see whether the database is up, enter `nbdb_ping`.

- `nbdb_admin [-start | -stop BMRDB]`

Starts or stops BMRDB without shutting down the NetBackup Scale-Out Relational Database server.

To see whether the BMRDB database is up, enter `nbdb_ping -dbn BMRDB`.

## The NetBackup database and clustered environments

The NetBackup database is supported in a clustered environment. Failover is included with the NetBackup server failover solution. The software is installed on all computers in the cluster.

The databases and the configuration files are installed in the following shared locations.

### Windows

NetBackup databases:

`shared_drive\VERITAS\NetBackupDB\data`

Configuration files:

`shared_drive\VERITAS\NetBackupDB\data\instance`

### UNIX

NetBackup databases:

`shared_drive/db/data`

Configuration files:

`/usr/opensv/var/global`

`shared_drive/db/data/instance`

## Post-installation tasks

The tasks that are described in the following topics are optional and can be performed after the initial installation:

- Change the database password.  
See [“Changing the NetBackup database password”](#) on page 688.
- Move the NetBackup databases (possibly to tune performance).  
See [“Moving a database after installation ”](#) on page 689.
- Recreate NBDB.  
See [“Creating the NBDB database manually”](#) on page 691.

## Commands and utilities for administering the NetBackup databases

---

**Note:** Using the database administration utilities to administer the NetBackup database can potentially break the consistency between the NetBackup catalog and the database. This loss of consistency can lead to loss of data. Only use these utilities and commands with assistance of Cohesity Technical Support.

---

The following utilities are available to administer the databases.

See [“Using the NetBackup Database Administration utility on Windows”](#) on page 693.

See [“Using the NetBackup Database Administration utility on UNIX”](#) on page 698.

Also see the following commands in the *NetBackup Commands Reference Guide*.

```
create_nbdb
```

```
nbdb_backup
```

```
nbdb_restore
```

```
nbdb_unload
```

## Changing the NetBackup database password

The database password is set to a randomly generated password upon installation. This password is used for NBDB and BMRDB and for all DBA and application accounts. You can use this procedure to change it to a known password.

The password is encrypted and stored in the `vxdbms.conf` file. The permissions for the `vxdbms.conf` file allow only a Windows administrator or a `root` user to read or write to it.

For requirements when NBAC is enabled, see the *NetBackup Security and Encryption Guide*.



### To change the database password

- 1 Log on to the server as a Windows Administrator or as `root`.
- 2 To change the password for the first time after installation, run the following command. The command updates the `vxdbms.conf` file with the new, encrypted string:

On Windows: `install_path\NetBackup\bin\nbdb_admin -dba new_password`

On UNIX: `/usr/openv/db/bin/nbdb_admin -dba new_password`

The password needs to be an ASCII string. Non-ASCII characters are not allowed in the password string.

- 3 To change a known password to a new password, you can either use the `nbdb_admin` command or the NetBackup Database Administration utility. You must know the current password to log into the NetBackup Database Administration utility.

See [“Using the NetBackup Database Administration utility on Windows”](#) on page 693.

See [“Using the NetBackup Database Administration utility on UNIX”](#) on page 698.

## Moving a database after installation

The NetBackup database (NBDB) and the NetBackup authorization database (NBAZDB), are created on the primary server by default. To improve performance, you can use the NetBackup database administration utilities or command-line options to change the location of the databases.

Note the following:

- If BMR is installed and you want to move its database, it must reside on the primary server.
- Due to performance issues, you can only move a database to another disk or volume. The disk or volume must be locally attached.  
NetBackup does not support saving the NetBackup database (NBDB, including EMM), NBAZDB, or the configuration files to a remote NFS share. CIFS is supported on some SAN storage and NAS storage.
- Run a catalog backup to back up NBDB and BMRDB both before and after moving the databases.

### Moving a NetBackup database on Windows

The following instructions describe how to use the database administration utility to move a database.

You can also use the following command:

```
install_path\Veritas\NetBackup\bin\nbdb_move.exe
```

You can run the `nbdb_move` command at any time because it does not drop the database and recreate it. Therefore all the data is preserved.

### To move a NetBackup database on Windows

- 1 Perform a catalog backup.
- 2 Shut down all NetBackup services by typing the following command:  

```
install_path\Veritas\NetBackup\bin\bpdown
```
- 3 Start the NetBackup Scale-Out Relational Database Manager service:  

```
install_path\Veritas\NetBackup\bin\bpup -e vrtsdbsvc_psql
```
- 4 Start the NetBackup Database Administration utility and enter the database logon password. Click **OK**.
- 5 From the **Database** list, select the database that you want to move.
- 6 Select the **Tools** tab.
- 7 Click **Move**.
- 8 Select **Move data to** and browse to the new location.
- 9 NetBackup does not require that the database directories are world-writable. Make sure that the new database directories (*data\_directory*) have appropriate permissions so that the directories are not world-writable.
- 10 Start all services by typing the following command:  

```
install_path\Veritas\NetBackup\bin\bpup
```
- 11 Perform a catalog backup.

## Moving a NetBackup database on UNIX

### To move a NetBackup database on UNIX

- 1 Perform a catalog backup.
- 2 Shut down all NetBackup daemons by typing the following command:  

```
/usr/opensv/netbackup/bin/bp.kill_all
```
- 3 Start the NetBackup Scale-Out Relational Database Manager daemon:  

```
/usr/opensv/netbackup/bin/nbdbms_start_stop start
```
- 4 Use one of the following methods to move the existing databases:

- Use the **Move Database** option in the NetBackup Database Administration utility (dbadm).

- Enter the following command:

```
/usr/opensv/db/bin/nbdb_move
-data data_directory
```

You can run the `nbdb_move` command at any time because it does not drop the database and recreate it. Thus, all data is preserved.

```
/usr/opensv/db/bin/nbdb_move -data data_directory
```

---

**Note:** NetBackup does not require that the database directories are world-writable. Make sure that the new database directories (`data_directory`) have appropriate permissions so that the directories are not world-writable.

---

- 5 Start all NetBackup daemons by typing the following command:

```
/usr/opensv/netbackup/bin/bp.start_all
```

- 6 Perform a catalog backup.

## Copying the NetBackup databases

A temporary backup of the NBDB, NBAZDB, and BMRDB databases can be made for extra protection before database administration activities such as moving or reorganizing the databases. Also, some customer support situations may require that you create a copy of the NetBackup database.

Use the NetBackup database administration utilities or the `nbdb_backup` command to make this kind of backup.

## Creating the NBDB database manually

The NBDB database is created automatically during NetBackup installation. However, it may be necessary during certain catalog recovery situations to create it manually by using the `create_nbdb` command.

---

**Caution:** Recreating the database manually is not recommended in most situations.

---

---

**Note:** If the NBDB database already exists, the `create_nbdb` command does not overwrite it. If you want to move the database, move it by using the `nbdb_move` command.

---

### To create the NBDB database manually on Windows

- 1 Shut down all NetBackup services by typing the following command:

```
install_path\Veritas\NetBackup\bin\bpdown
```

- 2 Start the NetBackup Scale-Out Relational Database Manager service with the following command:

```
install_path\Veritas\NetBackup\bin\bpup -e vrtsdbsvc_psql
```

- 3 Run the following command:

```
install_path\Veritas\NetBackup\bin\create_nbdb.exe
```

- 4 Start all NetBackup services by typing the following command:

```
install_path\Veritas\NetBackup\bin\bpup
```

- 5 The new NBDB database is empty and does not contain the EMM data that is loaded during a normal installation.

Make sure that you have the most current support for new devices before the data is repopulated. New devices are added approximately every 2 months.

- 6 Repopulate the EMM data by running the `tpext` utility. `tpext` updates the NetBackup database with new versions of device mappings and external attribute files.

```
install_path\Veritas\Volmgr\bin\tpext.exe
```

During regular installation, `tpext` is run automatically.

If the `create_nbdb` command is used to create a database manually, the `tpext` utility must also be run. `tpext` loads EMM data into the database.

### To create the NBDB database manually on UNIX

- 1 Shut down all NetBackup daemons by typing the following command:

```
/usr/opensv/netbackup/bin/bp.kill_all
```

- 2 Start the NetBackup Scale-Out Relational Database Manager service with the following command:

```
/usr/opensv/netbackup/bin/nbdbms_start_stop start
```

- 3 Run the following command:

```
/usr/opensv/db/bin/create_nbdb
```

- 4 Start all NetBackup daemons by typing the following command:

```
/usr/opensv/netbackup/bin/bp.start_all
```

- 5 The new `NBDB` database is empty and does not contain the `EMM` data that is loaded during a normal installation.

Make sure that you have the most current support for new devices before the data is repopulated. New devices are added approximately every 2 months.

- 6 Repopulate the `EMM` data by running the `tpext` utility. `tpext` updates the NetBackup database with new versions of device mappings and external attribute files.

```
/usr/opensv/volmgr/bin/tpext
```

During regular installation, `tpext` is run automatically.

If the `create_nbdb` command is used to create a database manually, the `tpext` utility must also be run. `tpext` loads `EMM` data into the database.

## Additional `create_nbdb` options

In addition to using the `create_nbdb` command to create the `NBDB` database, you also can use it to perform the following actions. In each command, *NB\_server\_name* matches the name in the following file: `postgresql.conf`

- Drop the existing `NBDB` database and recreate it in the default location:

```
create_nbdb -drop
```

On UNIX, the location of the current `NBDB` data directory is retrieved automatically from the `bp.conf` file.

- Drop the existing `NBDB` database and do not recreate it:

```
create_nbdb -drop_only
```

- Drop the existing `NBDB` database and recreate it in the *data* directory:

```
create_nbdb -drop -data data_directory
```

If the `NBDB` database was moved from the default location by using `nbdb_move`, use this command to recreate it in the same location. Specify `current_data_directory`. `BMRDB` must also be recreated. The `BMRDB` database must reside in the same location as the NetBackup database.

# Using the NetBackup Database Administration utility on Windows

The NetBackup administrator can use the Database Administration utility to configure the NetBackup databases and to monitor database operations. To use the utility, the administrator must have Administrator user privileges.

## General tab of the NetBackup Database Administration utility

The **General** tab contains information about database tablespaces. The tab contains tools to let the administrator reorganize fragmented database objects and validate and rebuild the database.

**Table 51-2** General tab options

| Option                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Refresh</b>        | Displays the most current information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Reorganize All</b> | This option defragments the tablespaces that are fragmented.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Validate</b>       | <p>This option performs a database validation on all of the database tablespaces in the selected database.</p> <ul style="list-style-type: none"> <li>Validates the indexes and keys on all of the tables in the database.</li> <li>Scans each table. For each row, a check is made that it exists in the appropriate indexes. The number of rows in the table must match the number of entries in the index.</li> <li>Ensures that every row that is referenced in each index exists in the corresponding table. For foreign key indexes, it also ensures that the corresponding row exists in the primary table.</li> </ul> <p>After a validation check runs, the Results screen lists each database object. Each error is listed next to the database object where it was found. The total number of errors are listed at the end of the list of database objects. If no errors were found, that is indicated.</p> <p>If any validation errors are reported, perform the following tasks:</p> <ul style="list-style-type: none"> <li>Shut down NetBackup (all daemons and services).</li> <li>Start only the NetBackup database server (vrtssdbsvc_psqli).</li> <li>Click <b>Validate</b> to repeat the validation check or use the <code>nbdb_admin.exe</code> command line utility.</li> </ul> <p>If validation errors persist, contact Cohesity Technical Support. The administrator may be asked to rebuild the database using the <b>Rebuild</b> option or the <code>nbdb_unload.exe</code> command line utility.</p> |
| <b>Rebuild</b>        | <p>This option unloads and reloads the database. A new database with all of the same options is built in its place.</p> <p>A Database Rebuild may be required if validation errors are reported when you use the <b>Validate</b> option.</p> <p><b>Note:</b> Before you rebuild the database, it is recommended that you create a copy of the database by performing a backup from the <b>Tools</b> tab.</p> <p>To rebuild the database temporarily suspends NetBackup operations and can take a long time depending on the database size.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

## About fragmentation

Table fragmentation can impede performance. When rows are not stored contiguously, or if rows are split into more than one page, performance decreases because these rows require additional page accesses.

When an update to a row causes it to grow beyond the originally allocated space, the row is split. The initial row location contains a pointer to another page where the entire row is stored. As more rows are stored on separate pages, more time is required to access the additional pages.

Reorganizing may also reduce the total number of pages that are used to store the table and its indexes. It may reduce the number of levels in an index tree. Note that the reorganization does not result in a reduction of the total size of the database.

The **Rebuild** option on the **General** tab completely rebuilds the database, eliminating any fragmentation, and free space. This option may result in a reduction of the total size of the database.

See [“Estimating catalog space requirements”](#) on page 668.

## Tools tab of the NetBackup Database Administration utility

The **Tools** tab of the NetBackup Database Administration utility contains a variety of tools to administer the selected database:

|                      |                                                                                                                  |
|----------------------|------------------------------------------------------------------------------------------------------------------|
| <b>Password</b>      | See <a href="#">“Changing the DBA password using the NetBackup Database Administration utility”</a> on page 695. |
| <b>Move Database</b> | See <a href="#">“Moving a NetBackup database”</a> on page 696.                                                   |
| <b>Unload</b>        | See <a href="#">“Exporting database schema and data”</a> on page 696.                                            |
| <b>Backup</b>        | See <a href="#">“Copying or backing up a database”</a> on page 697.                                              |
| <b>Restore</b>       | See <a href="#">“Restoring a database from a backup”</a> on page 697.                                            |

## Changing the DBA password using the NetBackup Database Administration utility

To change a known password to a new password, you can either use the `nbdb_admin` command or the NetBackup Database Administration utility.

### To change the DBA password from a known password to a new password

- 1 Select the **Tools** tab.
- 2 In the **Password** section, click **Change**.

- 3
- Enter the new password and confirm the new password. Changing the password changes it for both NBDB and BMRDB, if a BMR database is present.
- 4
- Enable **Create a backup file of your new DBA password** to keep track of the password.
- 5
- Click **OK**.  
  
The utility warns you that it is important to remember the password. You cannot recover information within the NetBackup database if the password is unavailable.
- 6
- Restart the database for the password change to take effect.

## Moving a NetBackup database

Use the NetBackup Database Administration utility to change the location of a database.

For full instructions on how to move a database, see the following topic.

See [“Moving a database after installation ”](#) on page 689.

## Exporting database schema and data

### To export database schema and data

- 1
- Select the **Tools** tab.
- 2
- In the **Unload** section, click **Export**.
- 3
- Browse to a destination directory.
- 4
- Select one or more of the following options:

|                 |                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Schema          | Unload only the database schema. The schema is unloaded as a file that is named <i>database.sql</i> in the named directory. For the NBDB database, the schema is unloaded as a file that is named <i>NBDB.sql</i> in the named directory. For other databases, a similar file is created. For example, for BMRDB the file is <i>BMRDB.sql</i> . For NBAZDB the file is <i>NBAZDB.sql</i> . |
| Schema and data | Unload both the database schema and the data. The data is unloaded as a set of files in comma-delimited format. One file is created for each database table.                                                                                                                                                                                                                               |

- 5
- Click **OK**.



## Copying or backing up a database

Use the NetBackup Database Administration utility to back up the database to a specified directory.

It is recommended that you create a backup copy of a database in the following situations:

Before you move the database.

See [“Moving a NetBackup database”](#) on page 696.

Before you rebuild the database.

See [“General tab of the NetBackup Database Administration utility”](#) on page 694.

---

**Note:** Using the NetBackup Database Administration utility to back up and restore the NetBackup database can potentially break the consistency between the NetBackup catalog and the database. This loss of consistency can lead to loss of data. Use the tool to back up and restore the NetBackup catalog only as a precautionary measure.

---

### To copy or back up a database

- 1 Start the NetBackup Database Administration utility and enter the database logon password. Click **OK**.
- 2 Select the **Tools** tab.
- 3 Click **Copy**.
- 4 Browse to a destination directory.

A copy of the database is made to this directory. This directory is also the location of the database that the **Restore** option uses.

---

**Note:** This backup is not a catalog backup, performed as part of regular NetBackup operations.

---

See [“Restoring a database from a backup”](#) on page 697.

- 5 Click **OK**.

## Restoring a database from a backup

Use the NetBackup Database Administration utility to restore a database from a backup copy.

The restore overwrites the current database. The database is shut down and restarted after the restore is completed.

A database restore causes NetBackup activity to be suspended, so do not perform a database restore while active backups or other restores run.

---

**Note:** Using the Database Administration utility to back up and restore the NetBackup database can potentially break the consistency between the NetBackup catalog and the database. This loss of consistency can lead to loss of data. Use the tool to back up and restore the NetBackup database only as a precautionary measure.

---

#### **To restore a database from a backup**

- 1 Start the NetBackup Database Administration utility and enter the database logon password. Click **OK**.
- 2 Select the **Tools** tab.
- 3 Click **Restore**.
- 4 Browse to the directory that contains the backup database.
- 5 Click **OK**.

## **Using the NetBackup Database Administration utility on UNIX**

The NetBackup Database Administration utility (`dbadm`) is a standalone application that is supported for NBDB and BMRDB. It is installed in the following location:

```
/usr/opensv/db/bin
```

To use the NetBackup Database Administration utility, you must be an administrator with root user privileges. When you start the NetBackup Database Administration utility, enter the DBA password. The password is set to a randomly generated password upon installation. Use the `nbdb_admin` command to change it to a known password if you have not done so already.

See [“Changing the NetBackup database password”](#) on page 688.

After you log on, the NetBackup Database Administration utility displays the following information about the current database:

**Table 51-3** NetBackup Database Administration utility properties

| Property          | Description                                                        |
|-------------------|--------------------------------------------------------------------|
| Selected database | The selected database: NBDB or BMRDB                               |
| Status            | The status of the selected database: UP or DOWN                    |
| Consistency       | The validation state of the selected database: OK, NOT_OK, or DOWN |

The initial screen also displays the following Database Administration main menu:

**Table 51-4** Database Administration main menu options

| Option                                      | Description                                                                                                                                                                                                                                                                                             |
|---------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Select/Restart Database and Change Password | This option displays the menu where you can select a database to start or stop, and to change database passwords.<br>See <a href="#">"Select/Restart Database and Change Password menu options"</a> on page 700.                                                                                        |
| Database Space Management                   | This option displays the menu where you can perform the following actions: <ul style="list-style-type: none"> <li>■ Generate a database space utilization report</li> <li>■ Reorganize fragmented database objects</li> </ul> See <a href="#">"Database Space Management menu options"</a> on page 700. |
| Transaction Log Management                  | This option is not supported.                                                                                                                                                                                                                                                                           |
| Database Validation Check and Rebuild       | This option displays the menu where you can validate and rebuild the selected database.<br>See <a href="#">"Database Validation Check and Rebuild menu options"</a> on page 701.                                                                                                                        |
| Move Database                               | This option displays the menu where you can change the location of the database tablespaces.<br>See <a href="#">"Move Database menu options"</a> on page 702.                                                                                                                                           |
| Unload Database                             | This option displays the menu where you can unload either the schema or the schema and data from the database.<br>See <a href="#">"Unload Database menu options"</a> on page 703.                                                                                                                       |
| Backup and Restore Database                 | This option displays the menu where you can choose the backup and restore options for the database.<br>See <a href="#">"Backup and Restore Database menu options"</a> on page 703.                                                                                                                      |
| Refresh Database Status                     | This option refreshes the Status and Consistency in the main menu.                                                                                                                                                                                                                                      |

## Select/Restart Database and Change Password menu options

The Select/Restart Database and Change Password menu contains the following options.

**Table 51-5** Select/Restart Database and Change Password options

| Option                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NBDB                    | Select NBDB and then view or modify the database using the other <code>dbadm</code> menu options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| BMRDB                   | Select BMRDB and then view or modify the database using the other <code>dbadm</code> menu options.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Start Selected Database | Starts the selected database.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Stop Selected Database  | Stops the selected database.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Change Password         | <p>Changes the password for the databases. The password is changed for both NBDB and BMRDB, if applicable. Restart the database for the password change to take effect.</p> <p>To log into the Database Administration utility, you must know the current DBA password.</p> <p>To change the password for the first time after installation, use the <code>nbdb_admin</code> command. The command updates the <code>vxdbms.conf</code> file with the new, encrypted string:</p> <p>See <a href="#">“Changing the NetBackup database password”</a> on page 688.</p> <p>To change a known password to a new password, you can either use the <code>nbdb_admin</code> command or the NetBackup Database Administration utility.</p> |

## Database Space Management menu options

You can use the Database Space Management option to perform the following functions:

- To report on database space utilization
- To reorganize fragmented database objects

**Table 51-6** Database Space and Memory Management options

| Option                   | Description                                                                                                                                                                                                                                                                                                                                |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Report on Database Space | <p>The report contains the tablespaces and the physical pathnames of the databases.</p> <p>For each tablespace, the report displays the name, the amount of free space in KBytes, and the file size in KBytes. The report also displays the amount of free space that remains on each of the file systems being used for the database.</p> |

**Table 51-6** Database Space and Memory Management options (*continued*)

| Option              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Database Reorganize | <p>Select this option to reorganize fragmented database tablespaces.</p> <p>These actions are performed from the Database Reorganize menu as follows:</p> <ul style="list-style-type: none"> <li>1) Defragment All<br/>This option automatically determines the tablespaces that are fragmented.</li> <li>2) Table Level Defragmentation<br/>This option generates a fragmentation report for each database table. For each table, the report includes the TABLE_NAME, number of ROWS, number of ROW_SEGMENTS, and SEGS_PER_ROW.<br/>In addition, a * displays in the ! column for an individual table if it will be automatically selected for reorganization by the Defragment All option.<br/>A row segment is all or part of one row that is contained on one page. A row may have one or more row segments. The ROW_SEGMENTS value indicates total number of row segments for the table. The SEGS_PER_ROW value shows the average number of segments per row, and indicates whether or not a table is fragmented.<br/>A SEGS_PER_ROW value of 1 is ideal, and any value more than 1 indicates a high degree of fragmentation. For example, a value of 1.5 means that half of the rows are partitioned.<br/>See <a href="#">“About fragmentation”</a> on page 695.</li> </ul> |

## Database Validation Check and Rebuild menu options

The Database Validation Check and Rebuild option lets you validate and rebuild the currently selected database.

**Table 51-7** Database Validation Check and Rebuild menu options

| Option              | Description                                                                               |
|---------------------|-------------------------------------------------------------------------------------------|
| Standard Validation | The standard type of validation is not supported. This option performs a full validation. |

**Table 51-7** Database Validation Check and Rebuild menu options (*continued*)

| Option           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Full Validation  | <p>This option performs a database validation on all of the database tablespaces in the selected database.</p> <ul style="list-style-type: none"> <li>Validates the indexes and keys on all of the tables in the database.</li> <li>Scans each table. For each row, a check is made that it exists in the appropriate indexes. The number of rows in the table must match the number of entries in the index.</li> <li>Ensures that every row that is referenced in each index exists in the corresponding table. For foreign key indexes, it also ensures that the corresponding row exists in the primary table.</li> </ul> <p><b>Note:</b> To perform a full database validation, shut down NetBackup and start only the database service.</p> <p>If any validation errors are reported, perform the following tasks:</p> <ul style="list-style-type: none"> <li>Shut down NetBackup (all daemons and services).</li> <li>Start only the NetBackup database server (vrtsdbsvc_psqli).</li> <li>Repeat the validation check using this tool or the <code>nbdb_admin</code> command line utility.</li> </ul> <p>If validation errors persist, contact Cohesity Technical Support. The administrator may be asked to rebuild the database using the Database Rebuild option or the <code>nbdb_unload.exe</code> command-line utility.</p> |
| Database Rebuild | <p>This option lets you rebuild the database. A Database Rebuild results in a complete unload and reload of the database. A new database with all of the same options is built in place. A Database Rebuild may be required if Database Validation errors are reported using the Standard or Full Validation options.</p> <p>During a Database Rebuild, all NetBackup operations are suspended.</p> <p>When you select this option, a message appears which recommends that you exit and create a backup using the Backup Database option before you rebuild the database. You then have the choice of whether to continue or not.</p> <p>See <a href="#">“Backup and Restore Database menu options”</a> on page 703.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## Move Database menu options

The Move Database menu option lets you change the location of a database. After you select Move Database, you are prompted for the directory name where you want to move the database.

For full instructions on how to move a database, see the following topic.

See [“Moving a database after installation ”](#) on page 689.

## Unload Database menu options

The Unload Database menu options let you unload either the schema or the schema and data from the `NBDB` or the `BMRDB` database.

A file is created that can be used to rebuild the database. If the data is also included in the unload, a set of data files in comma-delimited format is created.

The Unload Database menu contains the following options.

**Table 51-8** Unload Database menu options

| Option           | Description                                                                                                                                                                                                                                        |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Schema Only      | This option lets you unload only the database schema. For the <code>NBDB</code> database, the schema is unloaded as a file that is named <code>NBDB.sql</code> in the named directory. For <code>BMRDB</code> the file is <code>BMRDB.sql</code> . |
| Data and Schema  | This option lets you unload both the database schema and the data. The data is unloaded as a set of files. One file is created for each database table.                                                                                            |
| Change Directory | This option lets you change the directory location for the files that unload options (1) or (2) create.                                                                                                                                            |

## Backup and Restore Database menu options

The Backup and Restore Database menu options let you back up the NetBackup database to the specified directory. You can restore from a previously created backup.

It is recommended to create a backup copy of the databases in the following situations:

- Before you move the database.
- Before you rebuild the database.

---

**Note:** Using the NetBackup Database Administration utility to back up and restore the NetBackup database can potentially break the consistency between the NetBackup catalog and the database. This loss of consistency can lead to loss of data. Use the tool to back up and restore the NetBackup database only as a precautionary measure.

---

**Table 51-9** Backup and Restore Database menu options

| Option           | Description                                                                                                                                                                                                                                 |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Online Backup    | This option lets you make a copy of the databases while the databases are active. Other NetBackup activity is not suspended during this time.                                                                                               |
| Restore Backup   | This option lets you restore from a copy of the databases that was previously made with either options 1 or 2. The currently running databases are overwritten, and the database is shut down and restarted after the restore is completed. |
| Change Directory | This option lets you change the directory location for the databases that the backup options (1) or (2) create. This directory is the source of the databases for the restore option (3).                                                   |