

# Dell EMC NetWorker

Version 18.1

## CloudBoost 18.1 Integration Guide

302-004-434

REV 01

Copyright © 2016-2018 Dell Inc. or its subsidiaries. All rights reserved.

Published July, 2018

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS-IS." DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.  
Published in the USA.

Dell EMC  
Hopkinton, Massachusetts 01748-9103  
1-508-435-1000 In North America 1-866-464-7381  
[www.DellEMC.com](http://www.DellEMC.com)

# CONTENTS

	<b>PREFACE</b>	<b>7</b>
<b>Chapter 1</b>	<b>CloudBoost Integration</b>	<b>11</b>
	About the CloudBoost appliance.....	12
	CloudBoost appliances with NetWorker software.....	12
	CloudBoost requirements and considerations.....	13
	On-Prem CloudBoost Management Console browser requirement....	13
	Direct back up to the cloud with Linux and Windows clients.....	14
	Backup Amazon EC2 data to Amazon S3 storage.....	14
	Backup a Microsoft Azure virtual machine data to Azure blob	
	storage.....	15
	Cloud best practices.....	15
	Supported private clouds.....	17
	Supported public clouds.....	17
	Firewall port requirements.....	18
<b>Chapter 2</b>	<b>NetWorker with CloudBoost solution requirements</b>	<b>21</b>
	Solution requirements.....	22
	WAN requirements.....	22
	Virtual CloudBoost appliance requirements (VMware ESX).....	22
	Virtual CloudBoost appliance requirements (Amazon Web Services	
	EC2).....	23
	Virtual CloudBoost appliance requirements for Microsoft Azure....	24
	NetWorker client host requirements.....	24
	CloudBoost sizing and performance considerations.....	25
	Virtual CloudBoost appliance sizing.....	25
	CloudBoost metadata storage requirements.....	25
	End-to-end bottlenecks.....	26
	Minimum WAN requirements.....	26
	CloudBoost caching.....	26
	CloudBoost appliance cache sizing.....	26
<b>Chapter 3</b>	<b>Deploying the Virtual CloudBoost Appliance with VMware ESX</b>	<b>29</b>
	Virtual CloudBoost appliance installation (VMware ESX).....	30
	Deploy the virtual CloudBoost appliance.....	30
	Configure network settings for a CloudBoost appliance.....	31
<b>Chapter 4</b>	<b>Deploying the CloudBoost Appliance with Microsoft Azure</b>	<b>35</b>
	<b>Resource Manager</b>	<b>35</b>
	Integrate the CloudBoost appliance with Microsoft Azure.....	36
	Download the VHD files and JSON template.....	36
	Configure and deploy the CloudBoost appliance .....	36
	Use the Azure PowerShell to configure Microsoft Azure for the	
	CloudBoost appliance.....	36
	Use the Azure CLI to configure the CloudBoost appliance.....	37

	Start the CloudBoost virtual machine.....	43
	Set the FQDN.....	43
	Check the Microsoft Azure audit logs.....	44
	Verify network setup and status of the appliance.....	44
<b>Chapter 5</b>	<b>Deploying the CloudBoost Appliance in Amazon EC2</b>	<b>47</b>
	Deploy the virtual CloudBoost appliance in Amazon EC2.....	48
	Start the CloudBoost virtual machine.....	49
	Set the FQDN.....	49
	Verify network setup and status of the appliance.....	50
<b>Chapter 6</b>	<b>Configuring Network Settings for a CloudBoost Appliance</b>	<b>51</b>
	Network settings for a CloudBoost appliance.....	52
	Configure network settings for a CloudBoost appliance.....	52
	Configure CloudBoost to use a proxy.....	54
<b>Chapter 7</b>	<b>Configuring a New CloudBoost Appliance</b>	<b>57</b>
	Create and manage cloud profile for CloudBoost appliance.....	58
	Cloud profiles.....	58
	Create a cloud profile.....	58
	Edit a cloud profile.....	58
	Validate cloud storage credentials.....	59
	Enable remote client mounts.....	60
	Configure a new CloudBoost appliance.....	60
	Editing CloudBoost appliance configurations.....	62
<b>Chapter 8</b>	<b>Configuring NetWorker with a CloudBoost appliance</b>	<b>63</b>
	Configure a CloudBoost device by using an embedded NetWorker storage node.....	64
	Configure a CloudBoost device on an external storage node.....	67
	Troubleshoot CloudBoost device configuration issues.....	72
	Setting the configuration options for the CloudBoost SDK.....	72
	Improve clone performance.....	73
	Cannot retrieve the version of the CloudBoost appliance.....	74
	The selected CloudBoost appliance is unsupported for device type "CloudBoost".....	74
	Directory not found.....	74
	Unable to connect to the CloudBoost appliance: LOGON_FAILURE error.....	74
	Adding a CloudBoost 2.2.2 appliance fails with an error "unable to resolve".....	75
	Report information on cloud backup.....	75
<b>Chapter 9</b>	<b>Perform a CloudBoost Appliance Recovery</b>	<b>77</b>
	Recovering CloudBoost Appliance.....	78
<b>Chapter 10</b>	<b>Monitoring, Managing, and Supporting a CloudBoost Appliance</b>	<b>79</b>
	Monitoring CloudBoost.....	80
	Upgrade a CloudBoost appliance.....	80
	CloudBoost integration with EMC Secure Remote Services .....	81

Registering EMC Secure Remote Services.....	81
Installing the EMC Secure Remote Services gateway.....	81
Register CloudBoost with EMC Secure Remote Services.....	82
Increase the CloudBoost appliance site cache.....	83
Configuring average chunk size.....	83
Specifications for the chunk size setting.....	84

## CONTENTS

# PREFACE

As part of an effort to improve product lines, periodic revisions of software and hardware are released. Therefore, all versions of the software or hardware currently in use might not support some functions that are described in this document. The product release notes provide the most up-to-date information on product features.

If a product does not function correctly or does not function as described in this document, contact a technical support professional.

---

## Note

This document was accurate at publication time. To ensure that you are using the latest version of this document, go to the Support website at <https://support.emc.com>.

---

## Purpose

This document describes the integration between the NetWorker software and the CloudBoost appliance.

## Audience

This guide is part of the NetWorker documentation set, and is intended for use by system administrators who are responsible for setting up and maintaining backups on a network. Operators who monitor daily backups will also find this guide useful.

## Revision history

The following table presents the revision history of this document.

**Table 1** Document revision history

Revision	Date	Description
01	July 07, 2018	First release of this document for NetWorker 18.1.

## Related documentation

The following publications provide information about CloudBoost.

- *CloudBoost Release Notes*  
Contains information about new features and changes, fixed problems, known limitations, environment and system requirements for the latest release.

You may also find it helpful to refer to these NetWorker publications.

- *NetWorker Administration Guide*  
Describes how to configure and maintain the NetWorker software.
- *NetWorker Installation Guide*  
Provides information about how to install, uninstall, and update the NetWorker software for clients, storage nodes, and servers on all supported operating systems.

## Special notice conventions that are used in this document

The following conventions are used for special notices:

**NOTICE**

Identifies content that warns of potential business or data loss.

---

**Note**

Contains information that is incidental, but not essential, to the topic.

---

**Typographical conventions**

The following type style conventions are used in this document:

**Table 2** Style conventions

<b>Bold</b>	Used for interface elements that a user specifically selects or clicks, for example, names of buttons, fields, tab names, and menu paths. Also used for the name of a dialog box, page, pane, screen area with title, table label, and window.
<i>Italic</i>	Used for full titles of publications that are referenced in text.
Monospace	Used for: <ul style="list-style-type: none"> <li>• System code</li> <li>• System output, such as an error message or script</li> <li>• Pathnames, file names, file name extensions, prompts, and syntax</li> <li>• Commands and options</li> </ul>
<i>Monospace italic</i>	Used for variables.
<b>Monospace bold</b>	Used for user input.
[ ]	Square brackets enclose optional values.
	Vertical line indicates alternate selections. The vertical line means or for the alternate selections.
{ }	Braces enclose content that the user must specify, such as x, y, or z.
...	Ellipses indicate non-essential information that is omitted from the example.

---

You can use the following resources to find more information about this product, obtain support, and provide feedback.

**Where to find product documentation**

- <https://support.emc.com>
- <https://community.emc.com>

**Where to get support**

The Support website at <https://support.emc.com> provides access to licensing information, product documentation, advisories, and downloads, as well as how-to and troubleshooting information. This information may enable you to resolve a product issue before you contact Support.

To access a product specific Support page:

1. Go to <https://support.emc.com/products>.
2. In the **Find a Product by Name** box, type a product name, and then select the product from the list that appears.




3. Click .
4. (Optional) To add the product to **My Saved Products**, in the product specific page, click **Add to My Saved Products**.

### Knowledgebase

The Knowledgebase contains applicable solutions that you can search for by solution number, for example, 123456, or by keyword.

To search the Knowledgebase:

1. Go to <https://support.emc.com>.
2. Click **Advanced Search**.  
The screen refreshes and filter options appear.
3. In the **Search Support or Find Service Request by Number** box, type a solution number or keywords.
4. (Optional) To limit the search to specific products, type a product name in the **Scope by product** box, and then select the product from the list that appears.
5. In the **Scope by resource** list box, select **Knowledgebase**.  
The **Knowledgebase Advanced Search** panel appears.
6. (Optional) Specify other filters or advanced options.
7. Click .

### Live chat

To participate in a live interactive chat with a support agent:

1. Go to <https://support.emc.com>.
2. Click **Chat with Support**.

### Service requests

To obtain in-depth help from Licensing, submit a service request. To submit a service request:

1. Go to <https://support.emc.com>.
2. Click **Create a Service Request**.

---

### Note

To create a service request, you must have a valid support agreement. Contact a sales representative for details about obtaining a valid support agreement or with questions about an account. If you know the service request number, then directly enter the service request number in the `Service Request` field to get the valid details.

---

To review an open service request:

1. Go to <https://support.emc.com>.
2. Click **Manage service requests**.

### Online communities

Go to the Community Network at <https://community.emc.com> for peer contacts, conversations, and content on product support and solutions. Interactively engage online with customers, partners, and certified professionals for all products.

### How to provide feedback

Feedback helps to improve the accuracy, organization, and overall quality of publications. You can send feedback to [DPAD.Doc.Feedback@emc.com](mailto:DPAD.Doc.Feedback@emc.com).



# CHAPTER 1

## CloudBoost Integration

This section contains the following topics:

- [About the CloudBoost appliance](#)..... 12
- [CloudBoost appliances with NetWorker software](#)..... 12
- [Supported private clouds](#)..... 17
- [Supported public clouds](#)..... 17
- [Firewall port requirements](#)..... 18

## About the CloudBoost appliance

The CloudBoost appliance provides an integrated solution for existing supported backup environment by enabling the transfer of backups to public, hybrid, or private cloud storage. The CloudBoost appliance supports the following use cases: long-term retention to the cloud and backup to a private or public cloud.

CloudBoost decouples metadata from data. Encryption keys, metadata, and file system information are housed separately from the data, which removes a common bottleneck for cloud read/write operations. All advanced data services, such as chunking, encryption, inline deduplication, compression, and bulk data transfers are performed separately from metadata storage.

CloudBoost is available as a VMware virtual appliance and as a virtual appliance that resides in supported public clouds.

## CloudBoost appliances with NetWorker software

A NetWorker with CloudBoost environment can extend onsite data protection to the cloud through the following methods:

### Backup to the cloud

NetWorker with CloudBoost allows direct backup of on-premises clients to a range of private, public, and hybrid clouds. This solution allows clients to send backups directly to the object store with only the metadata being stored in the CloudBoost appliance. This distributed model where the CloudBoost appliance is not in the data path provides enhanced backup performance, scale, and client-side data reduction. The solution supports Client Direct backup to the cloud for Linux and Windows file systems and a broad range of enterprise applications. For applications that do not support Client Direct, use an external or embedded NetWorker Storage Node to perform backups directly to the cloud.

### Backup in public cloud

This solution allows protection of applications that run in public clouds such as AWS, AWS S3, Azure, and Azure blob storage. Similar to on-premises backups to the cloud, this solution allows Client Direct backup to the object store for applications that run in AWS EC2 and Azure compute instances. For applications that do not support Client Direct, use an external or embedded NetWorker Storage Node to perform backups directly to the cloud.

### Long-term retention or cloning to cloud:

This solution allows clone backups from a backup target to the cloud for long-term retention. The operational copy for backup and restore operations remains on the Data Domain host or any other backup target. The copy that is cloned to the cloud by NetWorker and CloudBoost is used for long-term retention of data.

This table details the module support matrix for CloudBoost.

**Table 3** CloudBoost module matrix

Module	Application	External or embedded Storage Node	Cloning	Client Direct	
				Linux x64	Microsoft Windows 64-bit
File System	Not applicable	Yes	Yes	Yes	Yes
Block Based Backup					
NetWorker Module for Databases and Applications (NMDA)	DB2	Yes	Yes	Yes	Yes
	Informix				
	SAP IQ				
	Lotus				Yes
	MySQL				Not applicable
	Oracle				Yes, except for AWS
	Sybase				Yes
NetWorker Module for Microsoft	Microsoft Exchange	Yes	Yes	Not applicable	Yes
	Microsoft Hyper-V				
	Microsoft SharePoint				
	Microsoft SQL				
NetWorker Module for SAP (NMSAP)	SAP HANA	Yes	Yes	Yes	Not applicable
	SAP with Oracle				Yes
NetWorker Snapshot Module (NSM)	Not applicable	Not applicable	Yes	Yes, only RHEL	Yes
VBA	VMware			Not applicable	Not applicable
vProxy	VMware				

## CloudBoost requirements and considerations

Before you can use the CloudBoost appliance to protect data in a NetWorker datazone, you must deploy a CloudBoost appliance in the environment.

Refer to the Online Software Compatibility Guide at <http://compatibilityguide.emc.com:8080/CompGuideApp/> for the complete list of supported products and versions.

## On-Prem CloudBoost Management Console browser requirement

The supported web browser for On-Prem CloudBoost Management Console is:

- Google Chrome
- Microsoft Internet Explorer
- Mozilla firefox

## Direct back up to the cloud with Linux and Windows clients

This use case is intended for when you have onsite infrastructure and want to use object storage for all backup workloads, including short-term backups for operational recovery and long-term backups for compliance.

The optional site cache reduces the impact of long-distance connectivity. The optional site cache also meets recovery-time objectives more quickly, because the most frequently used data is cached locally.

Direct backup to the cloud is recommended for the following use cases:

- Where a high bandwidth pipe to the object store is required.
- When backing up non-critical applications that can tolerate a higher SLA for backup and restore operations.

This figure displays Linux and Windows clients that are directly backed up to the cloud.

**Figure 1** Clients backed up directly to the cloud



For clients that cannot back up directly to the cloud, you can send backups through the CloudBoost appliance or an external NetWorker storage node to the cloud. However, routing through either the CloudBoost appliance or the external NetWorker storage node limits performance. Having the data path go directly from the client to the cloud is the most scalable, efficient, and optimal performance deployment model.

## Backup Amazon EC2 data to Amazon S3 storage

This use case is intended for workloads that run in public clouds and use S3 cloud object storage for backups, including short term backups for operational recovery and long term retention backups for compliance.

The following figure displays back up in Amazon EC2 to Amazon S3 storage.

**Figure 2** Back up in Amazon EC2 to Amazon S3 storage



The optional site cache service is unavailable when you deploy the CloudBoost appliance within Amazon EC2.

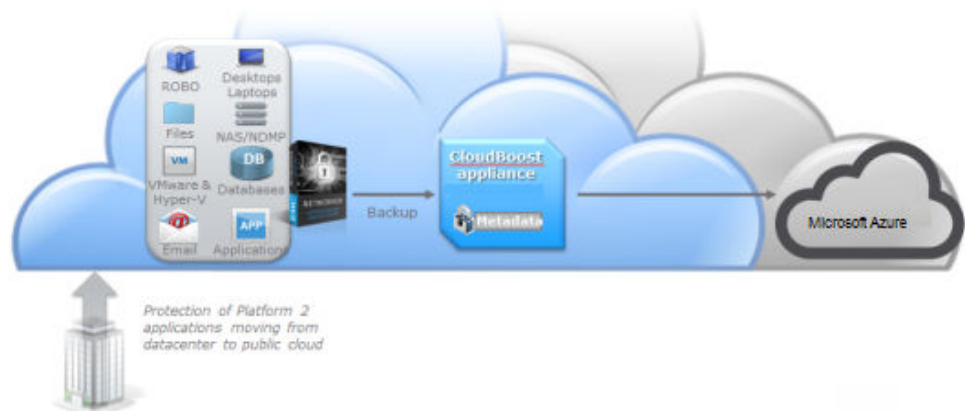
## Backup a Microsoft Azure virtual machine data to Azure blob storage

This use case is intended for workloads that run in the public cloud and use Microsoft Azure blob storage for backups, including short-term backups for operational recovery and long-term retention backups for compliance.

You use the same NetWorker tools to manage both onsite and cloud-based data protection processes.

The following figure illustrates backing up of Microsoft Azure virtual machine to Azure blob storage.

**Figure 3** Back up to Microsoft Azure



The optional site cache service is unavailable when you deploy the CloudBoost appliance within Microsoft Azure.

## Cloud best practices

Consider cloud best practices before you implement cloud backups.

## Backups to a cloud storage device

NetWorker backups are saved in media pools which can contain one or more backup devices. When a backup is triggered, the NetWorker server sends it to one of the unused devices in the media pool.

When creating a media pool that contains CloudBoost devices, do not include backup devices from other CloudBoost appliances or any other type of backup device. Each CloudBoost appliance has its own deduplication database and cannot deduplicate against backups that are sent to other devices.

## Concurrent backup and recovery operations

To support concurrent backup and recovery operations, you can create multiple NetWorker devices on a single CloudBoost appliance.

Each CloudBoost device type supports a minimum target session of 10 and a maximum session of 80. That means, each CloudBoost device supports a minimum of 10 concurrent streams and a maximum of 80 concurrent streams.

For example, to optimize performance you can mount the cloud volume on three cloud storage devices:

- One cloud storage device for backup (device CL1)
- One cloud storage device for recovery (device CL2)
- One cloud storage device for clone operations (device CL3)

You can create a maximum of 512 cloud devices per CloudBoost appliance, which is the maximum capacity of a NetWorker storage node. To optimize the backup and recovery performance, consider reducing the number of cloud devices per CloudBoost appliance. Each cloud device can handle a maximum of 80 concurrent streams.

## Network dependencies

Cloud backups depend on the network connection that accesses the cloud service. Any disruption in connectivity or a slowdown in network speed can adversely affect cloud backups or recoveries.

The CloudBoost appliance requires proper DNS name resolution and internet access.

Consider the following points before you set up the network for cloud backups:

- If the latency between the source and cloud object store is higher than 50 ms, backup and restore throughput from the object store might be impacted. NetWorker can sustain 100 ms on the metadata path. However, packet loss significantly impacts the backup success rate.
- If there is a high-latency link and some packet loss between the NetWorker server, client, and the CloudBoost appliance, set a high client-retry value for the backup so backups are re-tried.
- An increase of 5 ms latency in the data path (clients to the cloud object store, the CloudBoost appliance, or the cloud object store), has the following impacts:
  - For the initial full backup, throughput is two to two and a half times slower.
  - Consecutive backups are about 20 percent slower compared to a full backup.
- A higher-latency link and higher packet losses might result in significantly slower backup operations.



**Note**

It is recommended that latency between the NetWorker client and the cloud object store be limited to less than 50 ms and that packet loss be less than 1 percent.

## Supported private clouds

The following table lists the private clouds that CloudBoost appliances support.

**Table 4** Supported private clouds

Cloud provider	Information that is required by the CloudBoost appliance
Elastic Cloud Storage (ECS) Appliance	<ul style="list-style-type: none"> <li>ECS endpoint</li> <li>ECS access key ID</li> <li>ECS secret access key</li> </ul> <hr/> <p><b>Note</b></p> <p>Elastic Cloud Storage (ECS) Community version is not supported with CloudBoost.</p>
Generic OpenStack Swift	<ul style="list-style-type: none"> <li>Swift provider authentication endpoint</li> <li>Swift authentication type</li> <li>Region (optional)</li> <li>Swift credentials (specify the tenant name and the username separated by a colon, and then type the password)</li> <li>Swift secret key</li> </ul>

## Supported public clouds

The following table lists the public clouds that CloudBoost appliances support.

**Table 5** Supported public clouds

Cloud provider	Information that is required by the CloudBoost appliance
Amazon Web Services (S3)	<ul style="list-style-type: none"> <li>Storage region</li> <li>AWS access key ID</li> <li>AWS secret access key</li> </ul>
Microsoft Azure Storage (supports general purpose and blob storage accounts with hot and cool tiers)	<ul style="list-style-type: none"> <li>Azure account name</li> <li>Azure API key</li> </ul>

**Table 5** Supported public clouds (continued)

Cloud provider	Information that is required by the CloudBoost appliance
Virtustream Storage Cloud	<ul style="list-style-type: none"> <li>• Access Key ID</li> <li>• Secret Access Key</li> <li>• Endpoint URL</li> </ul>
Scality	<ul style="list-style-type: none"> <li>• Access Key ID</li> <li>• Secret Access Key</li> <li>• Endpoint URL</li> </ul>

## Firewall port requirements

As with all networked software solutions, adhering to best practices for security is encouraged to protect the deployment. If the ports in the following table are not configured before you configure the CloudBoost appliance, restart the CloudBoost appliance.

### Note

It is not recommended to route outbound http traffic from the CloudBoost appliance through a proxy because it can create a performance bottleneck. In environments where outbound http traffic is restricted, create an exception for the appliance in the firewall after you consult with the IT security team. To configure a proxy, see [Configure CloudBoost to use a proxy](#).

The following table outlines the firewall port requirements.

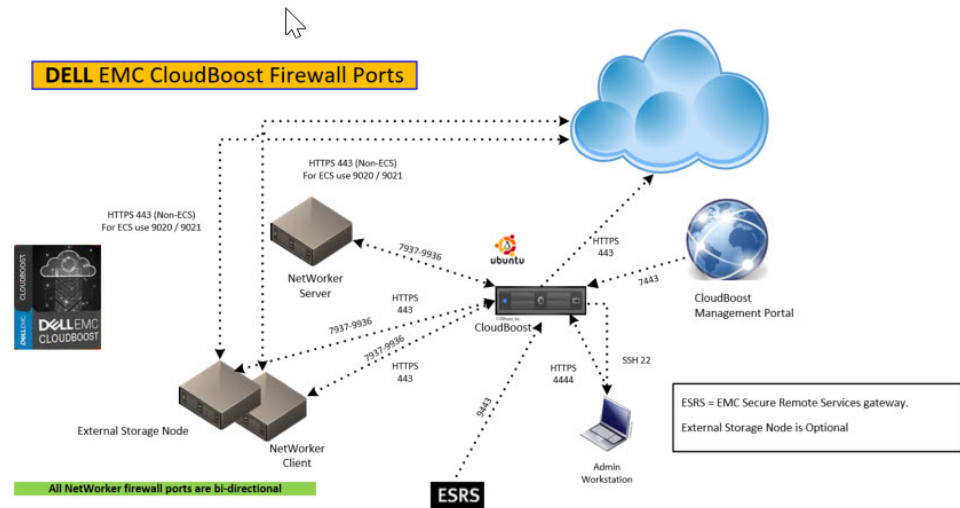
**Table 6** Firewall port requirements

Out	In	TCP port	Description
Administrator workstation	CloudBoost appliance	22	SSH for maintenance and troubleshooting
CloudBoost appliance	Cloud storage (public or private)	443	HTTPS to access object store (if supported)
CloudBoost appliance	On-Prem CloudBoost Management Console	7443	HTTPS to On-Prem CloudBoost Management Console.
NetWorker Server or NetWorker Client	CloudBoost appliance	7937-7942	The CloudBoost appliance has a pre-configured NetWorker SN. For a single CloudBoost device, a minimum of six ports must be opened on the CloudBoost appliance. The port range can be expanded based on the deployment type and the

**Table 6** Firewall port requirements (continued)

Out	In	TCP port	Description
			number of CloudBoost devices configured. The <i>NetWorker Security Configuration Guide</i> provides additional information on the NetWorker port requirements.
NetWorker client	<ul style="list-style-type: none"> <li>Cloud Storage (public or private)</li> <li>CloudBoost appliance for metadata.</li> </ul>	443	HTTPS to access object store (if supported)
CloudBoost appliance	EMC Secure Remote Services gateway	9443	Communication from the CloudBoost appliance to the Secure Remote Services gateway

**Figure 4** CloudBoost firewall ports



For information about firewall ports for any system that you deploy with CloudBoost, refer to the documentation for that system.

For information about NetWorker, refer to the *NetWorker Security Configuration Guide*.



# CHAPTER 2

## NetWorker with CloudBoost solution requirements

Before you begin the installation and configuration of the CloudBoost appliance, it is important that you understand all the requirements.

- [Solution requirements](#)..... 22
- [NetWorker client host requirements](#)..... 24
- [CloudBoost sizing and performance considerations](#)..... 25

## Solution requirements

This section outlines the solution requirements for the CloudBoost appliance in the following environments.

- VMware ESX
- Amazon Web Services (AWS) EC2
- Microsoft Azure

For more information on metadata store and site cache hard disk sizing, see [CloudBoost sizing and performance considerations](#) on page 25.

## WAN requirements

The following points provide the WAN requirements for the CloudBoost appliance.

- Greater than or equal to 100 Mb/s bandwidth
- Less than or equal to 100 ms RTT latency

## Virtual CloudBoost appliance requirements (VMware ESX)

The following section outlines the requirements and workflows that are supported for the VMware ESX virtual CloudBoost appliance.

### Minimum deployment virtual machine requirements for VMware ESX

This table outlines the requirements.

**Table 7** CloudBoost minimum requirements for VMware ESX

Workflow type	Site cache	CPU	Memory	OS	Metadata space	Site cache size
Client direct backup	No	4	16 GB	41 GB	100 GB	Not applicable
Backup/clone via CloudBoost appliance	No	8	32 GB			Not applicable
	Yes	16	64 GB			200 GB

### Large CloudBoost deployment virtual machine requirements for VMware ESX

This table outlines the requirements.

**Table 8** Large CloudBoost deployment requirements for VMware ESX

Workflow type	Site cache	CPU	Memory	OS	Metadata space	Site cache size
Client direct backup	No	8	32 GB	41 GB	100 GB	Not applicable
Backup/clone via CloudBoost appliance	No	16	64 GB	41 GB	Extendable up to 3 TB	Not applicable
	Yes	16	64 GB	41 GB	Extendable up to 3 TB	200 GB to 6 TB

---

**Note**

SSD is recommended for storage.

---

## Virtual CloudBoost appliance requirements (Amazon Web Services EC2)

The following section outlines the requirements and workflows that are supported for the Amazon Web Services (AWS) EC2 virtual CloudBoost appliance.

### Minimum deployment requirements for AWS EC2

This table outlines the minimum deployment requirements that apply to CloudBoost deployed on Amazon Web Services (AWS) EC2.

---

**Note**

Site cache is not supported for deployments on AWS EC2.

---

**Table 9** Minimum deployment requirements for AWS EC2

Workflow type	CPU	Memory	OS	Metadata space
Client direct backup	4	16 GB	41 GB	100 GB
Backup/Clone via CloudBoost appliance	8	32 GB		

Use the AWS EC2 instance, type m4.xlarge, for small-client direct installations.

---

**Note**

For smaller environments, you can choose an instance with unified compute and storage such as AWS EC2 m3.xlarge, which includes 4 vCPUs, 15 GB memory, and 2x40 GB SSD storage.

---

### Large deployment requirements for AWS EC2

This table outlines the requirements that apply to a large CloudBoost deployment on Amazon Web Services (AWS) EC2.

---

**Note**

Site cache is not supported for deployments on AWS EC2.

---

**Table 10** Large deployment requirements for AWS EC2

CPU	Memory	OS	Metadata space
8	32 GB	41 GB	100 GB metadata space is required per 200 TB of logical capacity

The primary metadata volume can be expanded to 3 TB to manage up to 6 PB of logical protected capacity.

Use the AWS EC2 instance type, m4.2xlarge, for Client Direct or for all installations of backup or clone operations through the CloudBoost appliance.

The CloudBoost appliance requires Amazon Elastic Block Store (AWS EBS) for the operating system disk and metadata database. Do not use the AWS EC2 instance

default storage volumes with the CloudBoost appliance. The AWS EC2 instance default storage volumes are ephemeral.

## Virtual CloudBoost appliance requirements for Microsoft Azure

The following section outlines the requirements and workflows that are supported for the Microsoft Azure virtual CloudBoost appliance.

### Minimum deployment virtual machine requirements for Microsoft Azure

This table outlines the minimum requirements are supported for Microsoft Azure.

**Table 11** Minimum requirements for Microsoft Azure

Workflow type	Site cache	CPU	Memory	OS	Metadata space
Client direct backup	No	4	16 GB	41 GB	100 GB metadata is required per 200 TB of logical capacity
Backup/clone via CloudBoost appliance	No	8	32 GB	41 GB	100 GB metadata is required per 200 TB of logical capacity
	No	16	64 GB	41 GB	100 GB metadata is required per 200 TB of logical capacity

## NetWorker client host requirements

This table outlines the NetWorker client host requirements for a virtual machine, public cloud, or private cloud.

**Table 12** NetWorker client host requirements

Workflow type	Operating System		CPU	Memory
	Linux x64	Microsoft Windows 64-bit		
NetWorker file system backup	Yes	Yes	2	4 GB
NetWorker Module backup	Yes	Yes	See <a href="#">CloudBoost appliance with NetWorker software</a> and the NetWorker Module documentation for specific requirements.	



# CloudBoost sizing and performance considerations

The following topics contain information about CloudBoost sizing, performance, and requirements.

## Virtual CloudBoost appliance sizing

For optimal performance, SSDs are recommended. Sizing for the virtual CloudBoost appliance depends on whether site caching is enabled.

- If site caching is enabled, 16 cores and 64 GB of memory is recommended.
- If site caching is not enabled, 8 cores with 32 GB of memory is recommended.

---

### Note

Site cache is unavailable on EC2 and Microsoft Azure deployments.

---

## CloudBoost metadata storage requirements

The amount of metadata storage that is required by a CloudBoost appliance depends on the average chunk size and data reduction ratio.

The virtual CloudBoost appliance requires a minimum of 100 GB of internal capacity for storing CloudBoost metadata. However, the amount of space that is provisioned for metadata directly affects the logical capacity addressable by the virtual CloudBoost appliance.

The ratio of metadata space to logical capacity ranges from 1:2000 to 1:500. For example, 100 GB of metadata allows the appliance to address 200 TB of logical capacity. To address the maximum logical capacity of 6 PB, 3 TB of metadata space is needed.

The virtual CloudBoost appliance assumes that the underlying storage is protected by using RAID or a similar technology. The virtual CloudBoost appliance does not provide protection against a failed virtual data disk.

Use the following formula to determine the CloudBoost metadata storage requirements:

```
Metadata storage = 100 GB (Reserved) + (512 / Data Reduction ratio) * (Logical Capacity in TB / Chunk Size in KB)
```

where:

- The default chunk size is 256 KB, which is the default.
- Deduplication ratio for the CloudBoost appliance is 2x–8x.

---

### Note

For best performance, use an SSD for metadata storage.

---

### Example 1 Examples

To address 6 PB of logical capacity with a dataset that has 4x data reduction, the CloudBoost appliance requires 3 TB of metadata storage.

**Example 1** Examples (continued)

To address 1 PB logical capacity with a dataset that has 4x data reduction, the CloudBoost appliance requires 500 GB for metadata.

## End-to-end bottlenecks

WAN bandwidth is expected to be the most common bottleneck. A properly resourced CloudBoost appliance can saturate a 1 Gb/s link with 30 ms RTT latency without hitting any limits in the virtual machine.

Object store ingest limits present another potential bottleneck. In some cases, we reach the objects/sec limit that can be sustained by a single logical container in the object store.

## Minimum WAN requirements

A minimum bandwidth of 100 Mb/s to the cloud with a maximum latency of less than 100 ms RTT is recommended for the CloudBoost solution. Extremely low bandwidth links might result in backup and restore timeouts.

## CloudBoost caching

The optional site cache allows quick backups over the LAN when data is moving slowly over the WAN.

Site caching enables faster backup and recovery for the objects that have been most recently written to or read from the cloud. This persistent cache is wiped and reused as needed during these processes.

If you intend to enable the site cache for a CloudBoost appliance, change the data disk size before you perform the initial configuration of the appliance at the CLI.

Use of the cache is advisable under the following circumstances:

- The connection to the object store is weak with low bandwidth and high latency, which is anything less than 200 Mb/s (25 MB/s) to the cloud store.
- You have no streaming workload or continuous backup.

---

**Note**

Site cache is available only on ESX deployments.

---

## CloudBoost appliance cache sizing

If you intend to enable the site cache for a CloudBoost appliance, change the data disk size before initial configuration of the CloudBoost appliance at the CLI.

Change the parameters in vCenter virtual machine configurations before you begin initial CloudBoost configuration at the CLI. It is recommended that you change these numbers to the appropriate levels based on the amount of data you plan to back up.

Consider the following information about CloudBoost appliance cache sizing:

- Without the site cache enabled, the ingestion rate for a CloudBoost appliance is measured at up to 100 MB/s.
- For the 32 TB virtual appliance, the site cache has a 50 MB/s ingestion rate.
- For all other appliances, the site cache has a 25 MB/s ingestion rate, which, depending on the workload, deduplication, and compression can improve.

- For simultaneous reading and writing, the underlying storage rate must be 25 MB/s.
- An average chunk size of 256 KB equals a few hundred IOPS.
- The total number of site cache disks must be a multiple of 2, up to 32. For example 1, 2, 4, 8, 16, or 32 disks.
- To optimize performance, configure each site cache VMDK on a different data store. For example, you can use four 500 GB VMDKs for a 2 TB site cache.
- You cannot increase or decrease the size of the cache by changing the existing data disk size in vCenter.
- The minimum size of the cache is 200 GB, which you can increase up to 6 TB on the virtual appliance by adding disks that match the size of the existing site cache disks.
- The cache is firewall-friendly, so you do not need to open multiple ports.

**Note**

If higher ingestion speeds occur when the CloudBoost appliance connects directly to the cloud store, do not use the cache. The backups might exceed the capacity of the cache. Do not use site cache if you are using the client direct workflow. The client direct workflow creates a bottleneck connecting to the site cache.

**Site cache minimum requirements**

This table outlines the minimum requirements for the site cache.

**Table 13** Site cache minimum requirements

Workflow type	Site cache	CPU	Memory	OS	Metadata	Site cache size
Backup/Clone	Yes	16	64 GB	41 GB	100 GB	200 GB



# CHAPTER 3

## Deploying the Virtual CloudBoost Appliance with VMware ESX

This section includes the following topics:

- [Virtual CloudBoost appliance installation \(VMware ESX\)](#)..... 30
- [Deploy the virtual CloudBoost appliance](#)..... 30
- [Configure network settings for a CloudBoost appliance](#)..... 31

## Virtual CloudBoost appliance installation (VMware ESX)

This section applies to installing the virtual CloudBoost appliance on VMware ESX. You must obtain the .OVA file from <https://support.emc.com> to install the virtual appliance.

The procedure for deploying the virtual appliance differs for each cloud provider:

- To deploy the virtual CloudBoost appliance on Amazon EC2, see [Deploying the CloudBoost Appliance in Amazon EC2](#) on page 47.
- To deploy the virtual CloudBoost appliance on Microsoft Azure, see [Deploying the CloudBoost Appliance with Microsoft Azure Resource Manager](#) on page 35.

## Deploy the virtual CloudBoost appliance

Deploy the virtual CloudBoost appliance in vSphere.

### Before you begin

- Determine the location of the .OVA file that you must download. This location could be a URL or a location that is accessible from the computer, such as a local hard drive or a network share.
- For the target data store, identify an available drive. If you choose to use a site cache, it is recommended that you use an SSD drive with at least 100 GB of available space plus the required space for the site cache.

### Procedure

1. In the vSphere client, click **File > Deploy OVF Template**.
2. Browse to the location of the OVA package, and then click **Next**.
3. Select the **Inventory Location** (the ESX cluster and host that runs the virtual machine), and then type the name of the virtual machine.
4. Select the data store for the VMDK files, and then click **Next**.

For optimal performance, select **Thick Provisioned Eager Zeroed** when you select the target data store. However, for testing purposes, the default 50 GB thin or thick provisioned storage is sufficient.

5. On the **Ready to Complete** page of the wizard, review the deployment settings.
6. Clear **Power on after deployment**, and then click **Finish**.
7. In the vSphere client inventory, right-click the virtual machine:
  - a. Click **Edit Settings**.
  - b. On the **Resources** tab, click **Memory**.
  - c. Ensure that **Reservation** is set to **a11**.
  - d. Click **OK**.

---

### Note

For performance, memory must be reserved and not shared.

---

8. Adjust the size of hard disk 2 (metadata store) and hard disk 3 (site cache) as required:
  - a. In the vSphere client inventory, right-click the virtual machine.
  - b. Click **Edit Settings**.
  - c. On the **Resources** tab, type the sizes of hard disk 2 and 3 in the required fields.
  - d. Click **OK**.

Any other hard disks that are added are used for the site cache.

The virtual CloudBoost appliance installs.

---

#### Note

Before you can finish configuring the virtual appliance in the [On-Prem CloudBoost Management Console](#), you must set the appliance's IP address and network through the CLI.

---

## Configure network settings for a CloudBoost appliance

### Procedure

1. Connect to the CloudBoost CLI.
2. Authenticate with the default password, `password`.
3. Set the new administrator password.
4. To see the current network configuration of the appliance, type the following command.

```
status
```

```
admin@mag-fs> status
Host Configuration:
  Hostname:      hostname
  Domain:       domain
  FQDN:         fqdn
Version Information:
  Version:      version identifier
  Internal Version: version identifier
  Revision:    revision identifier
Network Interfaces:
      name          mode          address
netmask  ----          ----          -----
-----
          eth0          dhcp          10.x.x.123      address

Network Routes:
      prefix          netmask          gateway
      ----          -----          -----
          default          0.0.0.0          10.x.x.1
          10.x.x.0          address          *
DNS Configuration
  DNS Servers:    10.x.x.91
Appliance Configuration
```

```
Status:      Not Configured
Endpoint:    NA
```

- If the IP address is dynamically assigned, then skip to [step 8](#). To statically set the IP address and netmask, type the following commands.

---

#### Note

If you have multiple networks, you must type the following commands for each network that is listed in the `status` command output.

---

- ```
net config interface static IP address netmask netmask
address
```

For example:

```
net config eth0 xx.x.xx.xxx netmask 0.0.0.0
```

- Manually add the gateway by typing these commands.

---

#### Note

If you have multiple networks, you must also add multiple routes to the gateways.

---

```
route add IP address netmask netmask address gw gateway
address
```

For example:

```
route add 0.0.0.0 netmask 0.0.0.0 gw xx.x.xx.x
```

- To manually set the DNS, type these commands:

```
dns set primary <primary IP address>
dns set secondary <secondary IP address>
dns set tertiary <tertiary IP address>
```

For example:

```
dns set primary 10.5.96.91
dns set secondary 10.5.96.92
dns set tertiary 10.5.96.93
```

- (Mandatory) To set the Fully Qualified Domain Name (FQDN), type the following command:

```
fqdn servername.yourcompanydomain
```

Consider the following:



- Custom *FQDN\_name* names are not supported.
- The *FQDN\_name* must be in lowercase.
- The *FQDN\_name* must not include the underscore character (\_).

For example:

```
fqdn cloudboost.example.com
```

---

**Note**

You must set the FQDN to access the On-Prem CloudBoost Management Console.

---

9. To verify the networking setup and to see the status of the appliance, type the following command:

```
status
```

**Results**

After you have verified the system's basic networking settings, configure CloudBoost by using the On-Prem CloudBoost Management Console.

---

**Note**

Other commands are available from the command line. To get help, type `help` or click the `?` icon.

---



# CHAPTER 4

## Deploying the CloudBoost Appliance with Microsoft Azure Resource Manager

Use the procedures in this section to deploy a CloudBoost appliance with Microsoft Azure Resource Manager (ARM).

- [Integrate the CloudBoost appliance with Microsoft Azure](#) ..... 36
- [Download the VHD files and JSON template](#)..... 36
- [Configure and deploy the CloudBoost appliance](#) ..... 36
- [Start the CloudBoost virtual machine](#).....43
- [Set the FQDN](#)..... 43
- [Check the Microsoft Azure audit logs](#)..... 44
- [Verify network setup and status of the appliance](#).....44

## Integrate the CloudBoost appliance with Microsoft Azure

Integrating the CloudBoost appliance with the Microsoft Azure cloud platform employs Microsoft Azure low-cost blob storage to provide deduplication at the source, which minimizes bandwidth and storage consumption.

Use the following procedures to deploy and integrate the CloudBoost appliance with Microsoft Azure.

## Download the VHD files and JSON template

Network bandwidth and large file sizes might cause the VHD and JSON file download to be time-consuming.

### Procedure

1. Download the CloudBoost appliance from Online Support at <https://support.emc.com>.
2. Extract the files from the downloaded Zip file.

---

### Note

7zip was used to zip these files. To download the files, you might need to install a 7zip compatible extractor.

---

The Zip file contains the following files:

- Virtual Root Hard Disk (VHD) Zip file:  
`./installer/target/azure/root.vhd`  
The root VHD file is approximately 3 GB. As a sparse file, the root VHD file is 50 GB.
- lvm VHD Zip file:  
`./installer/target/azure/lvm.vhd`  
The lvm VHD is approximately 50 MB. As a sparse file, the lvm VHD is 40 GB.
- JSON Zip file:  
`./management/mgt_console/resources/scripts/autobuild/azuredeploy.json`
- MD5 check sums for the root VHD and lvm VHD files:  
`./installer/target/azure/azure_discs_md5sum.txt`

## Configure and deploy the CloudBoost appliance

You can use either the Microsoft CLI or Azure PowerShell to configure and deploy the CloudBoost appliance.

## Use the Azure PowerShell to configure Microsoft Azure for the CloudBoost appliance

### Before you begin

The following software and permissions are required:

- Microsoft Azure PowerShell
- Microsoft Azure account
- Azure subscription

Refer to the Microsoft Azure documentation for installation and configuration details.

### Procedure

1. Start Microsoft Azure PowerShell.
2. Type the following command:

```
Add-AzureAccount
```

3. Type the following command:

```
Select-AzureSubscription -SubscriptionName subscription-name
```

where *subscription-name* is the Microsoft Azure subscription account.

4. Respond to the following prompts:

```
$storage_account_name = account_name
```

```
$rg_name = resourcegroup_name
```

```
$storage_account_key = account_key
```

where:

- *account\_name* is the name of the Microsoft Azure subscription account.
- *resourcegroup\_name* is name of the resource group.
- *account\_key* is the Microsoft Azure account key.

5. To deploy the CloudBoost virtual machine, use the following template:

```
New-AzureRmResourceGroupDeployment -<resourcegroup_name>  
$rg_name -<template_file> ./azure.json
```

where:

- *<resourcegroup\_name>* is name of the resource group.
- *<template\_file>* is the name of the Microsoft Azure template file.

## Use the Azure CLI to configure the CloudBoost appliance

Follow these steps to configure the CloudBoost appliance by using the Azure CLI.

## Configure Microsoft Azure for the CloudBoost appliance by using the Azure CLI

### Before you begin

The following software and permissions are required:

- Microsoft Azure account
- Azure CLI tools
- Azure utilities
- Go tools
- Git software

Refer to the Microsoft Azure documentation for installation and configuration details.

### Procedure

1. Authenticate Azure with the account.

- a. Open the Azure CLI, and then type the following command:

```
azure login
```

A device code appears with a link to the **Microsoft Azure Device Login** page.

- b. To open the **Microsoft Azure Device Login** page, click the link.
- c. In the **Code** field, type the device code.

2. In the Azure Resource Manager, perform the following steps:

- a. Create or use a storage account with the following properties:

- Deployment Model—Resource Manager
- Account Kind—General Purpose
- Replication—Select any type except ZRS.

- b. Set the deployment mode to Resource Mode. Type the following command:

```
azure config mode arm
```

3. In the Azure Resource Manager, perform the following steps:

- a. Create a resource group or use an existing resource group to deploy the CloudBoost Virtual Machine (VM).

This resource group is used in the steps that follow.

- b. Create a container in the storage account.

4. At the Azure CLI, retrieve the Azure storage account information.

- To list the storage account name, type, location, and resource group, type the following command:

```
azure storage account list
```

- To retrieve the Azure storage access key, type the following command:

```
azure storage account keys list <storage_account_name> -g
<resource_group>
```

5. Export or set the following environment variables that the Azure CLI will use. The procedure differs for Linux and Windows operating systems.

- On a Linux system, type the following commands:

```
export AZURE_STORAGE_ACCOUNT=<storage_account_name>
```

```
export AZURE_STORAGE_ACCESS_KEY=<primary_access_key>
```

---

#### Note

On a Linux OS, you can add these environment variables to the `.bashrc` file so you do not need to export the variables each time that you want to use them.

- On a Windows system, type the following commands:

```
set AZURE_STORAGE_ACCOUNT=<storage_account_name>
```

```
set AZURE_STORAGE_ACCESS_KEY=<primary_access_key>
```

## Methods for uploading the VHD files

You can use either the `azure-vhd-utils` Go tool or the Microsoft Azure CLI template to upload the VHD files to the Microsoft Azure storage account.

Uploading the files with the `azure-vhd-utils` Go tool is significantly faster than using the Microsoft Azure CLI template.

### Upload the .vhd files (azure-vhd-utils Go tool)

Upload the `.vhd` files by using the `azure-vhd-utils` Go tool that Microsoft provides.

#### Procedure

1. Install the latest version of the Go programming language software, which is available at the following link:  
<https://golang.org/dl/>
2. If required, install git.  
The git software is available at the following link: <https://git-scm.com/downloads>
3. Export or set the following `PATH` and `GOPATH` environment variables that the Azure CLI template will use. The procedure differs for Linux and Windows operating systems:

- On a Linux system, type the following commands:

```
export PATH=<Paths>;C:\Go\bin
```

```
export GOPATH=<GO_PATH>
```

---

#### Note

On Linux, you can add these environment variables to the `.bashrc` file, so you do not have to export the variables each time that you want to use them.

- On a Windows system, type the following commands:

```
set PATH=<Paths>;C:\Go\bin
```

```
set GOPATH=<GO_PATH>
```

The following site provides detailed information: <https://golang.org/doc/code.html#GOPATH>

4. Export or set the following environment variables that the Azure CLI template will use.

The procedure differs for Linux and Windows operating systems:

- On a Linux system, type the following commands:

```
export AZURE_STORAGE_ACCOUNT=<storage_account_name>
```

```
export AZURE_STORAGE_ACCESS_KEY=<storage_account_key>
```

where:

- `<storage_account_name>` is the name of the storage account where the blob is to be uploaded.
- `<storage_account_key>` is the storage account key information.

---

#### Note

On Linux, you can add these environment variables to the `.bashrc` file, so you do not have to export the variables each time that you want to use them.

- On a Windows system, type the following commands:

```
set AZURE_STORAGE_ACCOUNT=<storage_account_name>
```

```
set AZURE_STORAGE_ACCESS_KEY=<storage_account_key>
```

where:



- `<storage_account_name>` is the name of the storage account where the blob is to be uploaded.
- `<storage_account_key>` is the storage account key information.

5. To get the Microsoft Azure utilities, type the following command:

```
go get github.com/Microsoft/azure-vhd-utils
```

6. Upload the root VHD file by typing the following command on one line:

```
azure-vhd-utils upload --localvhddpath <root.vhd_path>
--stgaccountname <storage_account_name> --stgaccountkey
<storage_account_key> --containername
<container_name> --blobname <dest_blob_name>
```

where:

- `<root.vhd_path>` is the file path to the root VHD file.
- `<storage_account_name>` is the name of the storage account where the blob is to be uploaded.
- `<storage_account_key>` is the storage account key information.
- `<container_name>` is the destination location of the container in the storage account.
- `<dest_blob_name>` is the name of the blob in which to upload the VHD file.

7. Upload the LVM VHD file by typing the following command on one line:

```
azure-vhd-utils upload --localvhddpath <lvm.vhd_path>
--stgaccountname <storage_account_name> --stgaccountkey
<storage_account_key> --containername <container_name>
--blobname <dest_blob_name>
```

where:

- `<lvm.vhd_path>` is the file path to the LVM VHD file.
- `<storage_account_name>` is the name of the storage account where the blob is to be uploaded.
- `<storage_account_key>` is the storage account key information.
- `<container_name>` is the destination location of the container in the storage account.
- `<dest_blob_name>` is the name of the blob in which to upload the VHD file.

## Upload the .vhd files (Microsoft Azure CLI template)

The CLI template method can take several hours to upload the VHD files to the Microsoft Azure storage account. The azure-vhd-utils Go tool method can take less than an hour to upload the VHD files.

---

### Note

Depending on network bandwidth, uploading these files might be time consuming because of their large size.

---

### Procedure

1. Open the command prompt.
2. Upload the root and lvm VHDs to the storage account by typing the following command:

```
azure storage blob upload <image_to_upload> <container_name>
<blob_name>
```

where:

- *<image\_to\_upload>* is the root or lvm VHD file to upload.
- *<container\_name>* is the destination location of the container in the storage account.
- *<blob\_name>* is the name of the blob in which to upload the VHD file.

An output similar to the following appears:

```
$ azure storage blob upload CloudBoost-18.0.0.0-azure-root.vhd
vhds test -azure-root.vhd
info: Executing command storage blob upload
+ Checking blob test-azure-root.vhd in container vhds
+ Uploading CloudBoost-18.0.0-azure-root.vhd to
blob test-azure-root.vhd in container vhds
Percentage: 9.0% (3.69GB/41.00GB)
Average Speed: 1.86MB/S
Elapsed Time: 00:33:51 11
```

## Deploy the CloudBoost virtual machine with the JSON template

Deploy the CloudBoost virtual machine in Microsoft Azure by using the `azuredeploy.json` file that you downloaded from the Online Support at <https://support.emc.com>.

### Procedure

1. Deploy the CloudBoost virtual machine in Microsoft Azure. At the command prompt, type the following command:

```
azure group deployment create -n <deployment_name> -f
azuredeploy.json -g <resource_group>
```

where:

- *<deployment\_name>* is the name of the deployment. Use a unique name.
- *<resource\_group>* is the name of the resource group that deploys the virtual machine.

2. At the command prompt, type the following information:
  - For **vmName**, specify a unique name for the virtual machine.
  - For **storageAccountName**, specify the storage account name to associate with the virtual machine.
  - For **osDiskVhdUri**, type the destination URL of the following blob:  
`<URL_to_VHD_Azure_blob>/CloudBoost-2.2.2-azure-root.vhd`
  - For **metaDiskVhdUri**, type the destination URL of the following blob:  
`<URL_to_VHD_Azure_blob>/CloudBoost-2.2.2-azure-lvm.vhd`
3. To comply with the virtual machine size and location standards, adjust the JSON file. For detailed information about Azure virtual machine size, refer to the following Microsoft documentation:  
<https://docs.microsoft.com/en-us/azure/virtual-machines/virtual-machines-windows-sizes>

## Start the CloudBoost virtual machine

### Procedure

1. Log in to Microsoft Azure.
2. Select and open the CloudBoost appliance.
3. Connect to the CloudBoost appliance using either of the following methods:
  - Use an SSH client.
  - Log in through a browser.

The CloudBoost CLI appears.

4. Authenticate with the default password, `password`.
5. Set a new administrator password.

## Set the FQDN

### Procedure

1. Open the Azure CLI.
2. Log in to Microsoft Azure.
3. Select and open the CloudBoost appliance.  
 The CloudBoost CLI appears.
4. Authenticate with the new, previously set, administration password.
5. Set the FQDN. At the command prompt, typing the following command:

```
fqdn <FQDN_name>
```

Consider the following:

- Ensure that the *FQDN\_name* is the same name as the DNS Label Name that Microsoft Azure provides.

- Custom *FQDN\_name* names are not supported.
- The *FQDN\_name* must be in lowercase letters.
- The *FQDN\_name* must not include the underscore character (\_).

## Check the Microsoft Azure audit logs

To check the Microsoft Azure audit logs, at the command prompt, type the following command:

```
azure group log show <resource_group> deployment
```

where *<resource\_group>* is the name of the resource group that was used to deploy the virtual machine.

## Verify network setup and status of the appliance

### Procedure

1. To configure network settings for the CloudBoost appliance, see [Configuring Network Settings for a CloudBoost Appliance](#) on page 51.
2. To verify the network setup and see the status of the appliance, type the following command:

```
status
```

For example:

```
admin@mag-fs> status
Host Configuration:
  Hostname:      hostname
  Domain:       domain
  FQDN:         fqdn
Version Information:
  Version:      version identifier
  Internal Version: version identifier
  Revision:    revision identifier
Network Interfaces:
  name          mode          address
  ----          -
  eth0         dhcp         10.x.x.123   address
Network Routes:
  prefix        netmask      gateway
  -----
  default      0.0.0.0     10.x.x.1
  10.x.x.0     address     *
DNS Configuration
  DNS Servers:  10.x.x.91
Appliance Configuration
  Status:      Not Configured
  Endpoint:    NA
```

3. Configure the CloudBoost appliance, see [Configuring a New CloudBoost Appliance](#) on page 57.

4. To configure NetWorker with the CloudBoost appliance, see [Configuring NetWorker with a CloudBoost appliance](#) on page 63.



# CHAPTER 5

## Deploying the CloudBoost Appliance in Amazon EC2

This chapter includes the following topics:

- [Deploy the virtual CloudBoost appliance in Amazon EC2](#).....48
- [Start the CloudBoost virtual machine](#).....49
- [Set the FQDN](#)..... 49
- [Verify network setup and status of the appliance](#)..... 50

# Deploy the virtual CloudBoost appliance in Amazon EC2

## Before you begin

Ensure that Licensing has made the Amazon Machine Image (AMI) for the CloudBoost appliance available.

## Procedure

1. Log in to the Amazon Web Services console.
2. Click **EC2**.
3. Under **Images**, perform the following steps:
  - a. Select **AMIs**.
  - b. Click **Owned by me** and then, select **Private Images** from the list box.
  - c. Select **Private images**.
  - d. Click the CloudBoost image.
4. Under **Instance Type**, select an instance type that exceeds the following minimum requirements:
  - 4 CPUs
  - 16 GB of memory

For more information on system requirement, see [Virtual CloudBoost appliance requirements \(Amazon Web Services EC2\)](#) on page 23

It is recommended that you use `m4.xlarge`.

5. Under **Configure Instance Details**, perform the following steps:
  - a. Type the number of instances to create.
  - b. Select the network and submask.
  - c. Select **Auto-assign Public IP**.
6. Under **Add Storage** for the **Root** section:
  - a. Set **Size (GB)** to at least 41.

The default size for the added volume is 40 GB. Increase the default size for the added volume based on a 1:2000 ratio. This storage is used for the metadata store.
  - b. Add an EBS volume for metadata.
7. Under **Tag Instance**, define up to 10 keys to assist with AMI management and identification.
8. Under **Configure Security Group**:
  - a. To allow or deny public access, select an existing security group. Security groups are a set of firewall rules.
  - b. Consider the port requirements for the CloudBoost appliance.
9. Review information about the instance.
10. Select an existing key pair, or create a new key pair. You use this key to connect to the CloudBoost appliance.



11. Click **Launch instances**.
12. Before you can connect to the CloudBoost appliance, you must download the private key.

---

#### Note

Save the private key in a secure and accessible location. After the private key is created, you will be unable to download the private key again.

---

The CloudBoost appliance starts in Amazon EC2.

## Start the CloudBoost virtual machine

### Procedure

1. Log in to Amazon EC2.
2. Select the CloudBoost appliance, and then click **Connect**.
3. In the **Connect To Your Instance** wizard, connect with an SSH client.
4. Log in to the SSH terminal:

To log in to the SSH terminal as the Admin user, type the following command:

```
ssh -i admin@ AWS_FQDN_name_or_IP
```

Consider the following:

- Custom *FQDN\_name* names are not supported.
- The *FQDN\_name* must be in lowercase.
- The *FQDN\_name* must not include the underscore character (\_).

---

#### Note

It is best practice to keep the DHCP configuration options that Amazon supplied.

---

5. Select and open the CloudBoost appliance.  
The CloudBoost CLI appears.
6. Authenticate with the default password, `password`.
7. Set the new administrator password.

## Set the FQDN

### Procedure

1. Connect to the CloudBoost CLI.
2. Authenticate with the new administrator password that you configured in the previous step.
3. Set the FQDN by typing the following command:

```
fqdn <FQDN_name>
```

Consider the following:

- Custom *FQDN\_name* names are not supported.
- The *FQDN\_name* must be in lowercase.
- The *FQDN\_name* must not include the underscore character (\_).

## Verify network setup and status of the appliance

### Procedure

1. To configure network settings for the CloudBoost appliance, see [Configuring Network Settings for a CloudBoost Appliance](#) on page 51.
2. To verify the network setup and see the status of the appliance, type the following command:

```
status
```

For example:

```
admin@mag-fs> status
Host Configuration:
  Hostname:      hostname
  Domain:       domain
  FQDN:         fqdn
Version Information:
  Version:      version identifier
  Internal Version: version identifier
  Revision:    revision identifier
Network Interfaces:
  name          mode          address
  ----          -
  eth0         dhcp         10.x.x.123   address

Network Routes:
  prefix        netmask       gateway
  ----        -
  default      0.0.0.0      10.x.x.1
  10.x.x.0     address      *
```

```
DNS Configuration
  DNS Servers:  10.x.x.91
Appliance Configuration
  Status:      Not Configured
  Endpoint:    NA
```

3. Configure the CloudBoost appliance, see [Configuring a New CloudBoost Appliance](#) on page 57.
4. To configure NetWorker with the CloudBoost appliance, see [Configuring NetWorker with a CloudBoost appliance](#) on page 63.

# CHAPTER 6

## Configuring Network Settings for a CloudBoost Appliance

This section includes the following topics:

- [Network settings for a CloudBoost appliance](#) ..... 52
- [Configure network settings for a CloudBoost appliance](#)..... 52
- [Configure CloudBoost to use a proxy](#) ..... 54

## Network settings for a CloudBoost appliance

You must provide basic network settings information for a CloudBoost appliance at the CLI and complete initial configuration in the On-Prem CloudBoost Management Console.

### Note

In AWS, the CloudBoost AMI automatically uses the default VPC settings for the appliance IP address, DNS, and FQDN.

## Configure network settings for a CloudBoost appliance

### Procedure

1. Connect to the CloudBoost CLI.
2. Authenticate with the default password, `password`.
3. Set the new administrator password.
4. To see the current network configuration of the appliance, type the following command.

```
status
```

```
admin@mag-fs> status
Host Configuration:
  Hostname:      hostname
  Domain:       domain
  FQDN:         fqdn
Version Information:
  Version:      version identifier
  Internal Version: version identifier
  Revision:    revision identifier
Network Interfaces:
  name          mode          address
  ----          -
  eth0         dhcp         10.x.x.123   address
Network Routes:
  prefix        netmask      gateway
  ----        -
  default      0.0.0.0     10.x.x.1
  10.x.x.0     address     *
```

```
DNS Configuration
  DNS Servers:  10.x.x.91
Appliance Configuration
  Status:      Not Configured
  Endpoint:    NA
```

5. If the IP address is dynamically assigned, then skip to [step 8](#). To statically set the IP address and netmask, type the following commands.

---

**Note**

If you have multiple networks, you must type the following commands for each network that is listed in the `status` command output.

---

- ```
net config interface static IP address netmask netmask address
```

For example:

```
net config eth0 xx.x.xx.xxx netmask 0.0.0.0
```

**6. Manually add the gateway by typing these commands.**

---

**Note**

If you have multiple networks, you must also add multiple routes to the gateways.

---

```
route add IP address netmask netmask address gw gateway address
```

For example:

```
route add 0.0.0.0 netmask 0.0.0.0 gw xx.x.xx.x
```

**7. To manually set the DNS, type these commands:**

```
dns set primary <primary IP address>
dns set secondary <secondary IP address>
dns set tertiary <tertiary IP address>
```

For example:

```
dns set primary 10.5.96.91
dns set secondary 10.5.96.92
dns set tertiary 10.5.96.93
```

**8. (Mandatory) To set the Fully Qualified Domain Name (FQDN), type the following command:**

```
fqdn servername.yourcompanydomain
```

Consider the following:

- Custom `FQDN_name` names are not supported.
- The `FQDN_name` must be in lowercase.
- The `FQDN_name` must not include the underscore character (`_`).

For example:

```
fqdn cloudboost.example.com
```

---

**Note**

You must set the FQDN to access the On-Prem CloudBoost Management Console.

---

9. To verify the networking setup and to see the status of the appliance, type the following command:

```
status
```

**Results**

After you have verified the system's basic networking settings, configure CloudBoost by using the On-Prem CloudBoost Management Console.

---

**Note**

Other commands are available from the command line. To get help, type `help` or click the `?` icon.

---

## Configure CloudBoost to use a proxy

CloudBoost can be configured to use a proxy to communicate with the On-Prem CloudBoost Management Console.

Do not route outbound http traffic from the CloudBoost appliance through a proxy because it can create a performance bottleneck. In environments where outbound http traffic is restricted, create an exception for the appliance in the firewall after you consult with the IT security team.

In a private cloud environment, excludes are used when the CloudBoost appliance is in a proxy or isolated network. In this scenario, the IP of the private cloud is excluded. Each time that you run the `http-proxy` command, you override an earlier setting for excludes.

**Procedure**

1. Connect to the CloudBoost CLI.
2. Authenticate with the administrator password.
3. Check the status of the CloudBoost appliance by typing the following command:

```
status
```

If a proxy is set, this command displays and lists what is excluded.

4. To remove the current proxy configuration, type the following command:

```
http-proxy reset
```

Use this command to clear the last proxy setting and excludes

5. To add a proxy exclude, type the following command:

```
http-proxy <IP>:<port> exclude <prefix>/<mask>
```

where:

- `<IP>` is the IP address of the proxy server.

- *<port>* is the port number of the proxy server.
- *<prefix>* is the range of the source IP addresses for which you want to bypass the proxy.
- *<mask>* specifies the size of the range that is identified by the prefix.

For example, this command sets the proxy to 10.8.196.10:3128 and excludes 10.8.\*.\*. Note that (\*) is 0–255:

```
http-proxy 10.8.196.10:3128 exclude 10.8.0.0/16
```

6. To add exceptions for multiple exclusions, type the following command:

```
http-proxy <IP>:<Port> exclude <prefix>/<mask>,<prefix>/<mask>,...
```

where:

- *<IP>* is the IP address of the proxy server.
- *<port>* is the port number of the proxy server.
- *<prefix>* is the range of the source IP addresses for which you want to bypass the proxy.
- *<mask>* specifies the size of the range that is identified by the prefix.

For example, this command sets the proxy to 10.8.196.10:3128 and excludes 192.8.\*.\* and 100.9.22.24. Note that (\*) is 0–255:

```
http-proxy 10.8.196.10:3128 exclude
192.0.0.0/8,10.10.0.0/16,10.9.22.24/32
```





# CHAPTER 7

## Configuring a New CloudBoost Appliance

After you install a CloudBoost appliance and configure it at the CLI and complete its configuration in the On-Prem CloudBoost Management Console. You must create a cloud profile for the storage provider that the appliance will use before you can complete its configuration.

- [Create and manage cloud profile for CloudBoost appliance](#).....58
- [Validate cloud storage credentials](#).....59
- [Enable remote client mounts](#).....60
- [Configure a new CloudBoost appliance](#)..... 60
- [Editing CloudBoost appliance configurations](#)..... 62

# Create and manage cloud profile for CloudBoost appliance

The following topics describe how to create and manage cloud profiles for CloudBoost appliances.

## Cloud profiles

Before you configure a CloudBoost appliance in the On-Prem CloudBoost Management Console, create a cloud profile for the storage that the appliance will use.

## Create a cloud profile

Before you configure a CloudBoost appliance in the On-Prem CloudBoost Management Console, create a cloud profile for the storage the appliance uses.

### Before you begin

Obtain the necessary credentials for the cloud provider that you intend to use.

### Procedure

1. Use a web browser and sign in to the On-Prem CloudBoost Management Console as the administrator. Type the On-Prem CloudBoost Management Console address in the following format:

```
https://<FQDN of the appliance>:7443
```

---

### Note

- Only the administrator can log into the On-Prem CloudBoost Management Console.
  - The default username is admin and use the password that you updated during deployment of CloudBoost.
- 

2. In the left menu, click **Cloud Profiles**.  
The **Cloud Profiles** page opens.
3. To create a cloud profile, click **New Cloud Profile**.
  - a. In the **Display Name** field, type the name for this cloud profile.
  - b. In the **Cloud Storage Provider** field, select the cloud provider.
  - c. In the fields that appear for the selected cloud provider, provide the credentials and any additional information that is required to access this particular cloud object store.
  - d. Click **Save**.

## Edit a cloud profile

### Procedure

1. From the **On-Prem CloudBoost Management Console**, open the **Cloud Profiles** page.
2. To change information for an existing cloud profile, click **Edit**.

- a. On the **Edit a Cloud Profile** page, change any fields necessary.
- b. Click **Save**.

## Validate cloud storage credentials

Use the cloud storage credential validator, sometimes referred to as the *blobstore validator (BSV)*, to validate the cloud storage credentials that you intend to use with the CloudBoost appliance.

### Before you begin

- Configure the CloudBoost appliance with a valid cloud storage provider.
- The NTP must be configured.

---

#### Note

If the date and time of the ESX host, CloudBoost virtual machine and object store are out of sync, then the validation fails. You can sync these to a specific time in the same zone or to an NTP.

---

### Procedure

1. Connect to the CloudBoost CLI.
2. To see a list of valid cloud profiles, type the following command:

```
diagnostics bsv-cli "--cloud_profile_id="
```

---

#### Note

The quotation marks are required.

The result should be similar to the following output, but with a list of available cloud profiles.

```
Can't find cloud profile with ID . Possible values are:
1   VSC Virtustream Storage Cloud standard Storage
% Can't find cloud profile with ID
```

3. To validate the storage credentials against the listed profiles, type the following command:

```
diagnostics bsv-cli "--cloud_profile_id=1"
```

---

#### Note

The quotation marks are required.

The *#* represents the cloud profile to validate as listed in the result of Step 2.

The result should indicate that various BSV CLI commands are being validated.

```
Running BSV CLI with java options: -Djclouds.trust-all-
certs=true -Djclouds.s3.virtual-host-buckets=false
```

```
Running BSV CLI with arguments: --provider=atmos --
identity=05832cb9d39a40af96aaafd4a406aa6f/
A8581914817a4a8c264d
--credential=fXPGEkxSCo9Zt6QHLtg05I/axjc=
--endpoint=https://api.atmosonline.com validate
...
continues to validate
```

## Enable remote client mounts

When you enable remote client mounts, you create the password that you must use for the task. Share the admin username and new password with anyone who needs to remotely mount with the client.

### Procedure

1. Open a CLI window for the appliance.
2. Log in with the admin username and password.
3. Type the following command, where the value for *password* is the password that you created.

```
remote-mount-password enable password
```

## Configure a new CloudBoost appliance

After you provide basic network information at the CLI, you must use a web browser to finish configuration in the On-Prem CloudBoost Management Console. You can change certain configuration information for an appliance after the initial configuration.

### Before you begin

Define a cloud profile for use with this appliance.

### Procedure

1. Use a web browser and sign in to the On-Prem CloudBoost Management Console as the administrator. Type the On-Prem CloudBoost Management Console address in the following format:

```
https://<FQDN of the appliance>:7443
```

---

### Note

Only the administrator can log into the On-Prem CloudBoost Management Console.

---

2. In the left menu, click **CloudBoost Appliances**.  
The **Appliances** page opens.
3. Click **CloudBoost Appliances**, click the appliance that you want to configure.
4. To change the display name for this appliance from the default FQDN that you set in the CLI, in the **Name** field, type the new display name.
5. To prevent this appliance from using a site cache, clear **Enable Site Cache**.

---

**Note**

The site cache settings cannot be changed after configuration.

---

6. To minimize clock drift:
  - a. Select **Enable NTP**.
  - b. Type the hostname or IP address for at least one NTP server.

Consider the following:

- Use the NTP details of the host system on which NetWorker server is installed.
- You can specify multiple servers to provide redundancy in case one or more time servers are unavailable. However, the NTP server validation is done only for one NTP server.
- If the NTP is not enabled, then the CloudBoost appliance inherits the time zone from the host ESXi server.

7. To set the frequency of backups, select a schedule for **Backup Frequency**.
- 

**Note**

The backups referred to here are for the system state of the appliance and for the stored metadata. This is not a reference to any backup software integration.

---

8. To use asymmetric encryption keys:
    - a. Select **Enable backup encryption with asymmetric keys**.
    - b. Refer to the displayed instructions to help you create the private and public encryption keys. Only this method of asymmetric key creation is supported for the CloudBoost appliance.
    - c. From the output file that is a result of step b, copy the entire public key and paste it into the text box below the instructions on the **Configure** tab.
    - d. From the output file that is a result of step b, copy the entire private key and paste it somewhere safe. If you created a pass phrase, copy that as well.
- 

**Note**

You must safely store the private key and pass phrase. They must be provided to decrypt a recovered backup. Appliances that are backed up using the public key that is provided on the **Configure** tab cannot be recovered without the private key and pass phrase.

---

9. Review the selections and click **Save** to save these settings for the appliance. CloudBoost is configured.
10. Download the recovery metadata and the private key.

---

**Note**

- a. After editing the initial configurations, you must safely store the recovery files. They must be provided during a Disaster Recovery. Appliances that are backed up cannot be recovered without the private key and recovery metadata.
  - b. After you acknowledge, the recovery metadata file and private key will be purged from the CloudBoost. If you do not acknowledge, the file will be available for download in the **Pending Action**.
- 

## Editing CloudBoost appliance configurations

You can change certain configuration information for an appliance after the initial configuration.

**Procedure**

1. Use a web browser and sign in to the On-Prem CloudBoost Management Console as the administrator. Type the On-Prem CloudBoost Management Console address in the following format:

```
https://<FQDN of the appliance>:7443
```

---

**Note**

Only the administrator can log into the On-Prem CloudBoost Management Console.

---

2. In the left menu, click **CloudBoost Appliances**.

The **Appliances** page opens.

3. Click **Edit**
- 

**Note**

You can edit only the NTP server field and the backup frequency.

---

4. Review the changes and click **Save** to save these settings for the appliance. CloudBoost is configured.
  5. Download the recovery metadata and the private key.
- 

**Note**

- a. You must safely store the recovery metadata and the private key. They must be provided during a Disaster Recovery. Appliances that are backed up cannot be recovered without the private key and recovery metadata.
  - b. After you acknowledge, the recovery metadata file and private key is purged from the CloudBoost. If you do not acknowledge, the file will be available for download in the **Pending Action**.
-

# CHAPTER 8

## Configuring NetWorker with a CloudBoost appliance

This chapter applies to configuring NetWorker with a CloudBoost appliance by using backup to the cloud. Cloud-based data protection occurs over a TCP/IP network.

The CloudBoost appliance has an embedded storage node which can be used with the NetWorker server. However, it is recommended that you install the NetWorker storage node on a separate Linux or Windows server and do not use the embedded NetWorker storage node on the CloudBoost appliance.

- [Configure a CloudBoost device by using an embedded NetWorker storage node](#) .....64
- [Configure a CloudBoost device on an external storage node](#)..... 67
- [Troubleshoot CloudBoost device configuration issues](#).....72
- [Report information on cloud backup](#).....75

## Configure a CloudBoost device by using an embedded NetWorker storage node



To configure a CloudBoost device to receive backup or clone data, perform the following steps.

You do not require an external storage node when you back up Linux and Windows hosts by using Client Direct.

### Procedure

1. In the **NetWorker Administration** interface
  - a. Click **View > Diagnostic Mode**.
  - b. Ensure that NetWorker **Client Direct** is selected.  
Client Direct backups are enabled by default.
2. If remote client mounts are not configured, on the CloudBoost appliance, enable remote client mounts, and then define a password for the *remotebackup* user account:
  - a. Connect to the CloudBoost appliance with the admin account.
  - b. Type the following command:

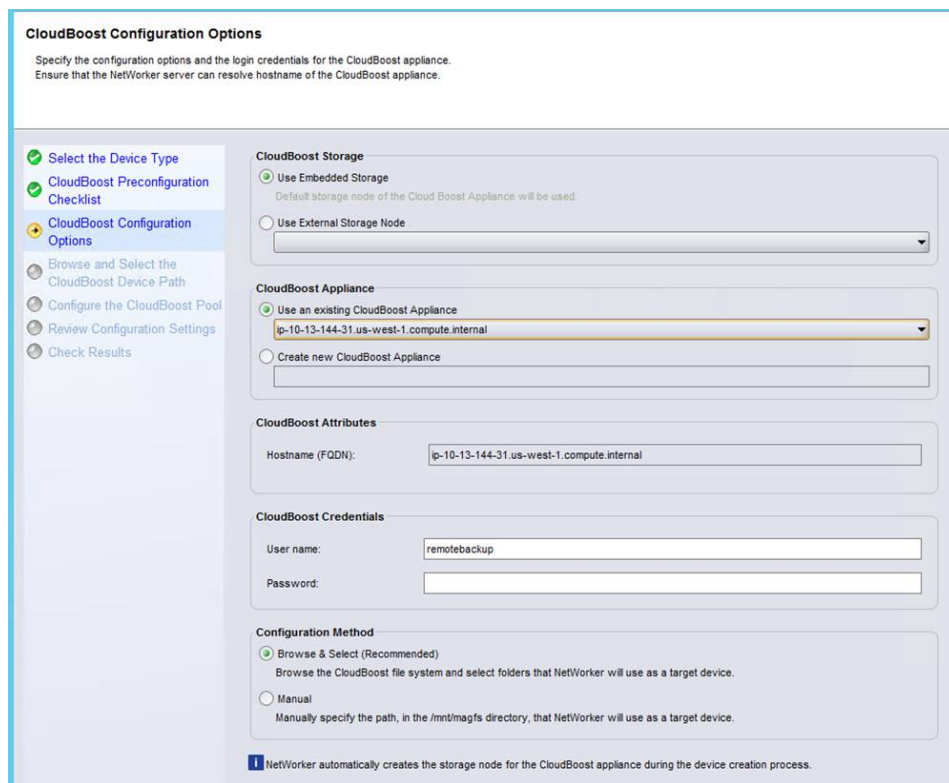
```
remote-mount-password enable password
```

where *password* is the new password for the *remotebackup* user.
3. Log in to the NMC GUI as an administrator of the NetWorker server.
4. On the taskbar, click the **Enterprise** icon .
5. In the navigation tree, highlight a host:
  - a. Right-click **NetWorker**.
  - b. Select **Launch Application**. The **NetWorker Administration** window appears.
6. On the taskbar, click the **Devices** button .
7. In the expanded left navigation pane:
  - a. Right-click **CloudBoost Appliances**.
  - b. Select **New Device Wizard**.
8. On the **Select the Device Type** page, select **CloudBoost**, and then click **Next**.
9. Review the **CloudBoost Preconfiguration Checklist** page, and then click **Next**.
10. On the **CloudBoost Configuration Options** page, perform the following tasks:
  - a. In the **CloudBoost Storage** group box:
    - a. Select **Use Embedded Storage Node**.
    - b. Select an embedded storage node.
  - b. In the **CloudBoost appliance** group box, select one of the following options:
    - To use a CloudBoost appliance that you have previously configured on the NetWorker server, select **Use an existing CloudBoost appliance**.



- To create a CloudBoost appliance, select **Create a new CloudBoost appliance** and then type a descriptive name.
- c. In the **Hostname (FQDN)** field, type the Fully Qualified Domain Name (FQDN) of the CloudBoost appliance.
  - d. In the **Username** field, type `remotebackup`.
  - e. In the **Password** field, type the password for the `remotebackup` account, which you defined on the CloudBoost appliance by using the `remote-mount` command.
  - f. In the **Configuration Method** group box, select the file system on the CloudBoost appliance that NetWorker uses as the target data device:

**Figure 5** Device Configuration Wizard: CloudBoost Configuration Options



- a. Select **Browse & Select**.  
The **Browse and Select the CloudBoost Device Path** window appears.
- b. In the `/mnt/magfs/base` directory, create a folder with write access enabled. Use a unique name.  
For example: `/mnt/magfs/base/CBO1`

Configure the CloudBoost appliance to cloud profile, the CloudBoost appliance creates a share on `/mnt/magfs/base`. The NetWorker software requires that each CloudBoost device has a unique, customer named, folder.

**Note**

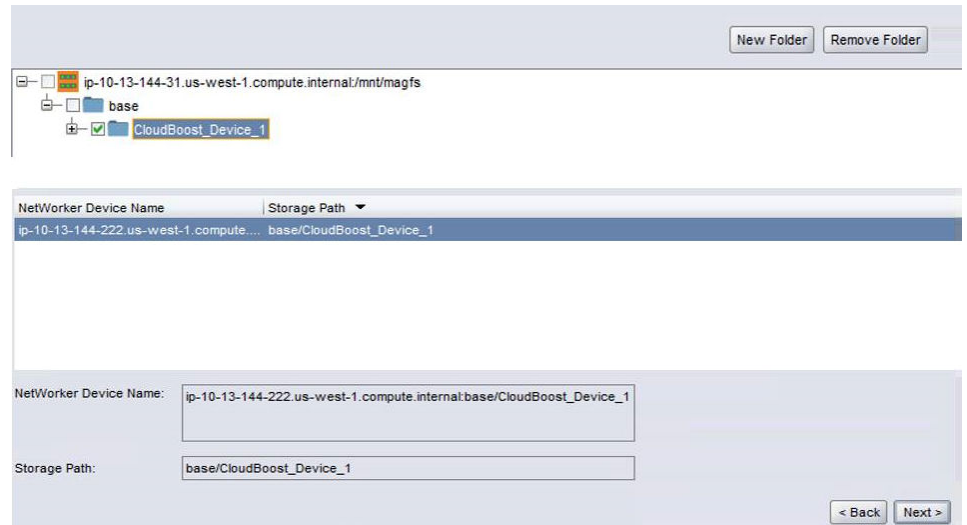
You cannot create folders outside of the `/mnt/magfs/base` directory.

c. Select the folder that you just created.

NetWorker validates the CloudBoost username and password that you specified. NetWorker also updates the **NetWorker Device Name** and **Storage Path** fields with the information.

g. Click **Next**.

**Figure 6** Device Configuration Wizard: Browse and Select the CloudBoost Device Path



11. On the **CloudBoost Pool Configuration** page, perform the following steps:

**Note**

For detailed information about NetWorker media pools, refer to the *NetWorker Administration Guide*.

- a. Ensure that **Configure Media Pools for devices** option is selected.
- b. In the **Devices** table, select the NetWorker device for the CloudBoost appliance.
- c. In the **Pool Type** box, depending on the use case, select one of the following:
  - **Backup**
  - **Backup Clone**
- d. In the **Pool** box:
  - To use a new pool, select **Create and use a new Pool**, and then type a pool name.
  - To use an existing pool, select **Use an existing Pool** and choose a pool that contains at least one CloudBoost device.

**Note**

The pool that you select cannot contain other device types such as AFTD and DD Boost devices.

- e. Ensure that **Label and Mount device after creation** is selected.
- f. Click **Next**.

**Figure 7** Device Configuration Wizard: Configure the CloudBoost Pool

Storage Path	NetWorker Device Name	Pool Type	Pool	Label
base/CloudBoost_De...	ip-10-13-144-222.us-west-1.com...	Backup	ip101314431uswest1...	✓

Backup  
 Backup Clone

Create and use a new Pool  
 Use an existing Pool

ip101314431uswest1computeinternal

Pool1

Label and Mount device after creation

12. On the **Review the Device Configuration** page:
  - a. Review the settings.
  - b. Click **Configure**.
13. On the **Check results** page:
  - a. Review whether the devices were successfully configured or if any messages appeared.
  - b. Click **Finish**.
  - c. To change any of the settings, click **Back** to the correct wizard page.


#### After you finish


If the Device Configuration wizard does not create a CloudBoost device, manually create an embedded storage node that corresponds to the embedded storage node that you created with the wizard. After the nrsnmd daemon starts on the CloudBoost appliance, create a CloudBoost device.

## Configure a CloudBoost device on an external storage node

For large scale deployments, the recommendation is to use an external storage node.

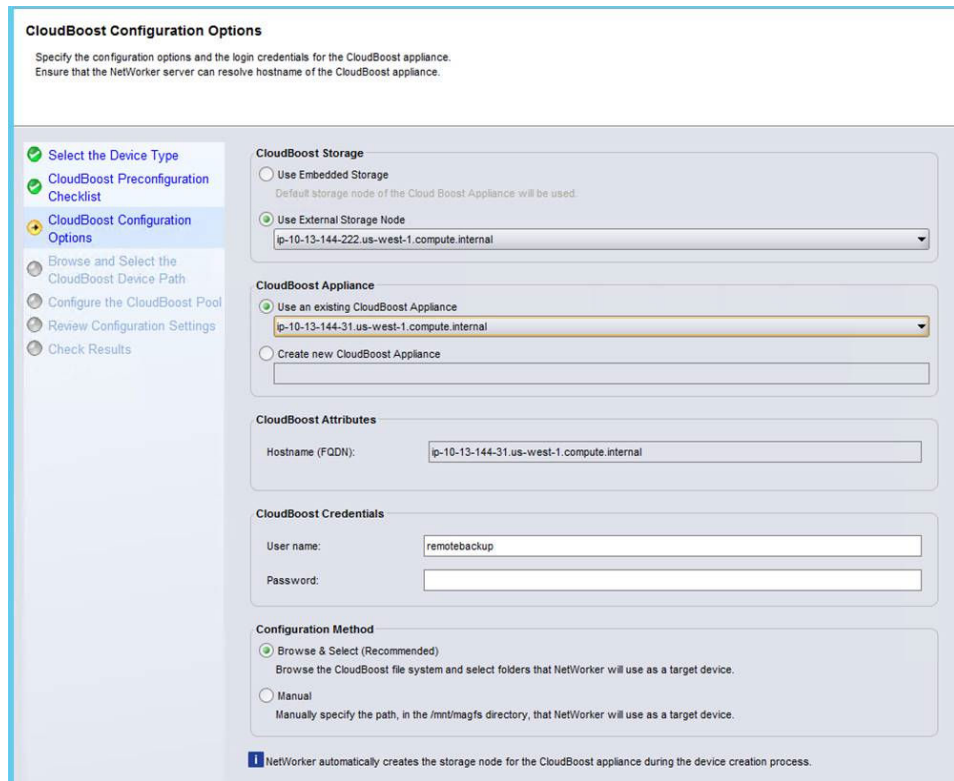
#### Procedure

1. In the NetWorker Administration interface, click **View > Diagnostic Mode**, and ensure that NetWorker **Client Direct** is selected.  
Client Direct backups are enabled by default.
2. Log in to the NMC GUI as an administrator of the NetWorker server.
3. On the taskbar, click the **Enterprise** icon .
4. In the navigation tree, highlight a host:
  - a. Right-click **NetWorker**.
  - b. Select **Launch Application**. The **NetWorker Administration** window appears.

5. Create a storage node:
  - a. From the navigation tree, right-click **Storage Nodes** and select **New**.  
The **Create Storage Node** windows appear with the **General** tab displayed.
  - b. Set the **Identity** attributes:
    - a. In the **Name** field, specify the hostname of the NetWorker storage node.
    - b. In the **Type of Storage Node** field, select **SCSI**.
  - c. In the **Status** attributes, review or set the storage node status:
    - **Storage node is configured** indicates that a device has been configured on this storage node.
    - **Enabled** indicates that the storage node is available for use:
      - **Yes** indicates an available state.
      - **No** indicates a service or disabled state. New device operations cannot begin and existing device operations might be canceled.
    - **Ready** indicates that the storage node is ready to accept device operations.
6. On the taskbar, click the **Devices** button .
7. Expand **Devices** in the left navigation pane:
  - a. Right-click the CloudBoost device.
  - b. Select **New Device Wizard**.The **Device Configuration Wizard** window appears.
8. On the **Select the Device Type** page:
  - a. Select **CloudBoost**.
  - b. Click **Next**.
9. Review the **CloudBoost Preconfiguration Checklist** page, and then click **Next**.
10. On the **CloudBoost Configuration Options** page, perform the following tasks:
  - a. In the **CloudBoost Storage** group box:
    - a. Select **Use External Storage Node**.
    - b. Select an external storage node.
  - b. In the **CloudBoost Appliance** group box, select one of the following options:
    - To use a CloudBoost appliance that you have previously configured on the NetWorker server, select **Use an existing CloudBoost appliance**.
    - To create a CloudBoost appliance, select **Create a new CloudBoost appliance** and specify a descriptive name.
  - c. In the **Hostname (FQDN)** field, specify the Fully Qualified Domain Name (FQDN) of the CloudBoost appliance.
  - d. In the **Username** field, type `remotebackup`.
  - e. In the **Password** field, type the password for the `remotebackup` account, which you defined on the CloudBoost appliance by using the `remote-mount` command.

- f. In the **Configuration Method** group box, select **Browse & Select**.  
The **Browse and Select the CloudBoost Device Path** page appears.

**Figure 8** Device Configuration Wizard: CloudBoost Configuration Options



- 11. On the **Browse and Select the CloudBoost Device Path** page, select the file system on the CloudBoost appliance that NetWorker uses as the target data device:
  - a. Select **New Folder**.
  - b. Create a folder in the `/mnt/magfs/base` directory. Use a unique name.

For example: `/mnt/magfs/base/CB01`

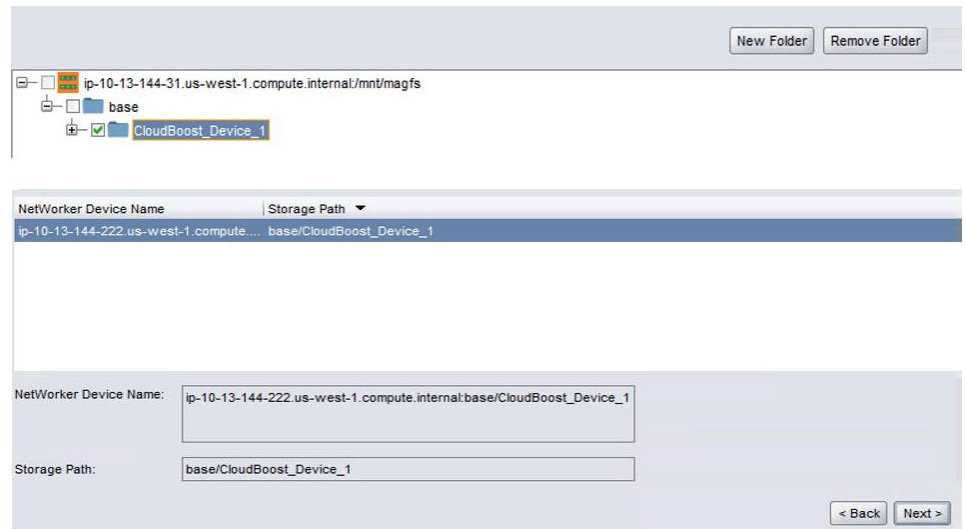
Configure the CloudBoost appliance to cloud profile, the CloudBoost appliance creates a share on `/mnt/magfs/base`. The NetWorker software requires that each CloudBoost device has a unique, customer named, folder.

**Note**

You cannot create folders outside of the `/mnt/magfs/base` directory.

- c. Select the folder that you just created.
- d. Ensure that the storage path that you specify exists in a subfolder in the `/mnt/magfs/base` directory with write access enabled.  
  
NetWorker updates the **NetWorker Device Name** and **Storage Path** fields with the required information.
- e. Click **Next**.

**Figure 9** Device Configuration Wizard: Browse and Select the CloudBoost Device Path



12. On the **Configure the CloudBoost Pool** page, perform the following steps:

**Note**

For detailed information about NetWorker media pools, refer to the *NetWorker Administration Guide*.

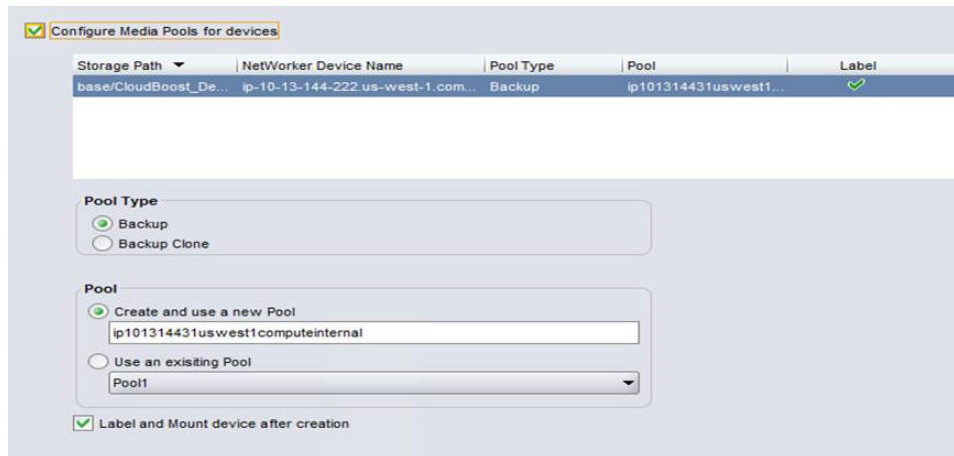
- a. Ensure that the **Configure Media Pools for devices** is selected.
- b. In the **Devices** table, select the NetWorker device for the CloudBoost appliance.
- c. In the **Pool Type** box, depending on the deployment, select either of the following:
  - **Backup**
  - **Backup Clone**
- d. In the **Pool** box, perform either of the following steps:
  - To use a new pool, select **Create and use a new Pool** and specify a pool name.
  - To use an existing pool, select **Use an existing Pool** and select a pool that contains at least one CloudBoost device.

**Note**

The pool that you select cannot contain other device types such as AFTD and DD Boost devices.

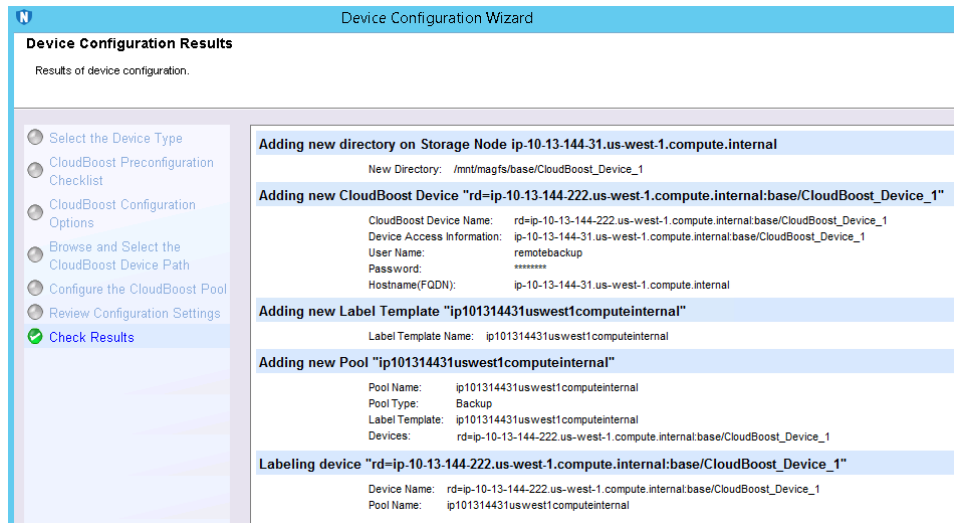
- e. Leave **Label and Mount device after creation** selected.
- f. Click **Next**.

**Figure 10** Device Configuration Wizard: Configure the CloudBoost Pool



13. Review the configuration settings, and then click **Next**.
14. On the **Check Results** page, check and verify that the device configuration was successful.

**Figure 11** Device Configuration Wizard: Check Results Page




**Results**

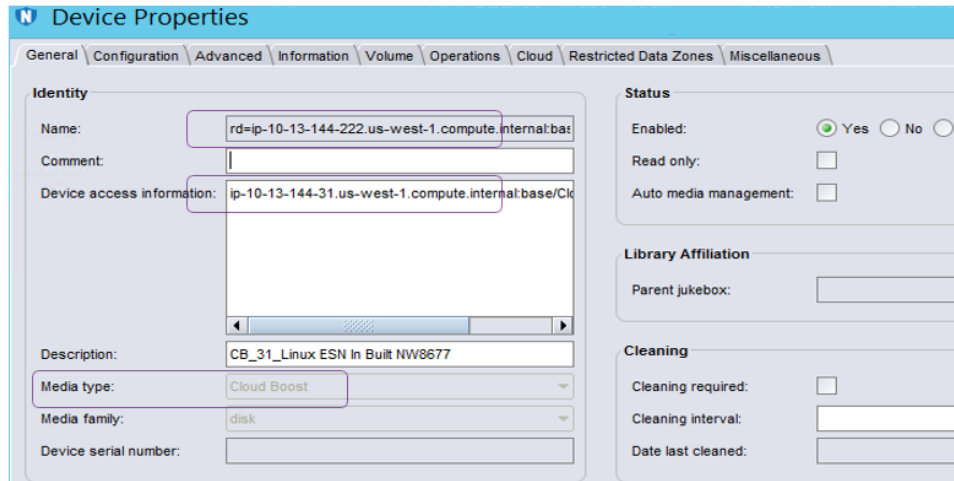
After the CloudBoost device has been configured with the NetWorker external storage node, the following information displays:

- The **Name** field points to the NetWorker external storage node.
- The **Device access information** field points to the CloudBoost appliance.
- The media type is specified as **CloudBoost**.

**Figure 12** Device information



**Figure 13** Device properties



## Troubleshoot CloudBoost device configuration issues

This section provides information about the error messages that might appear when you configure a NetWorker device for the CloudBoost appliance.

### Setting the configuration options for the CloudBoost SDK

You can set the configuration option for the CloudBoost SDK by creating the `nsrccbconfig` file to pass specific set of configuration option to CloudBoost library.

#### Procedure

1. Perform one of the following action sequences depending on the operating system:

Operating system	Procedure
Windows	<ol style="list-style-type: none"> <li>a. Open Notepad, and create the <code>nsrccbconfig</code> file with the configuration options.</li> <li>b. Save the file in <code>&lt;installation Directory&gt;\Emc Networker\nsr\debug\</code> directory without an extension.</li> </ol>
Linux	<ol style="list-style-type: none"> <li>a. Open the terminal.</li> <li>b. Create <code>nsrccbconfig</code> file without an extension in the <code>/nsr/debug/</code> directory and update the file with the configuration options.</li> </ol>

You can add one or more configuration options from the following table.



**Table 14** Configuration options

Configuration options	Description
v (lower case only)	Provides the capability of different log levels that ranges from 0-5. Greater the number, more the debug logs printed in /nsr/logs/cloudboost/*
maxLogFiles	Controls the number of log files that are retained in /nsr/logs/cloudboost/ folder. The default value is 10. The maximum value is 50.
max_log_size	Limits the size of log files. When the value reaches the maximum file size, the files is zipped and the log continues in a new file. The default size of log file is 10 MB.
curlConnectTimeout	This is the timeout value for establishing connection to object store. This can be set to a higher value (in msec) if there are connection failures with an error message "BAD_NETWORK_PATH or Transfer error". The default value is 20000 msec.
cacheLocation	It is the directory where CloudBoost SDK stores the blob cache. You must create the directory incase if it does not exist. The default location is /nsr/logs/cloudboost/ folder.
log_dir	It is the directory where the CloudBoost SDK writes log files. If it is not present, then you must create it. The default location is /nsr/logs/cloudboost/ folder.

The following example creates a nsrconfig file in /nsr/debug/nsrconfig location.

```
[root@ip-xx-xx-xx-xx /]# cat /nsr/debug/nsrconfig
v=5
maxLogFiles=15
max_log_size=100
curlConnectTimeout=30000
```

## Improve clone performance

For a new installation of NetWorker, the **Disable (RPS) Clone** attribute is selected by default. For NetWorker 9.0.1 and later, the **Disable (RPS) Clone** attribute is cleared by default. If you are upgrading to NetWorker 18.1, the **Disable (RPS) Clone** attribute keeps the previously assigned setting.

To improve clone performance, consider the following facts:

- For concurrent cloning workloads, select the **Disable (RPS) Clone** attribute.
- Ensure that the number of source and destination volumes match. This step avoids contention between the source and destination volumes and reduces the chances of clone failure.

To clear or select the **Disable (RPS) Clone** attribute, perform the following steps.

### Procedure

1. Open the **Administration** window.

2. Right-click the NetWorker server name in the left pane.
3. Select **Properties**.
4. In the **Configuration** tab, clear or select the **Disable RPS Clone** attribute.

## Cannot retrieve the version of the CloudBoost appliance

This error message appears when the NetWorker server cannot determine the version of NetWorker that is running on the CloudBoost appliance.

To resolve this issue, contact Technical Support for access to the CLI and ensure that the following criteria are met.

- The NetWorker daemon `nsrexecd` is started on the appliance:

1. Log in to an SSH terminal with the Maginatics user account.
2. Type the following command:

```
ps -ef | grep nsr
```

3. Confirm that the `nsrexecd` process appears.

4. If the `nsrexecd` daemon does not appear, type the following command to start the `nsrexecd` process on the appliance:

```
sudo service networker start
```

- Forward and reverse name resolution is correctly configured for the CloudBoost appliance and the NetWorker server host.

## The selected CloudBoost appliance is unsupported for device type "CloudBoost"

If the NetWorker server and the CloudBoost appliance cannot communicate with each other, this error message appears.

The host and DNS entries were not updated between the CloudBoost appliance and the NetWorker.

## Directory not found

This error appears when the NetWorker server cannot access the file system on the CloudBoost appliance.

To resolve this issue, ensure that the `/mnt/magfs` folder exists on the CloudBoost appliance and is mounted.

---

### Note

For NetWorker 9.1.1 and later, the `/mnt/magfs/base` folder is validated when a CloudBoost device is created. You cannot create a device in a folder other than `/mnt/magfs/base`.

---

## Unable to connect to the CloudBoost appliance: LOGON\_FAILURE error

This error appears in the following scenarios:

- The password that is specified for the `remotebackup` user is incorrect.
- The `remotebackup` user does not exist on the CloudBoost appliance.

## Adding a CloudBoost 2.2.2 appliance fails with an error “unable to resolve”

Perform the following steps to add a CloudBoost appliance 2.2.2 in the NetWorker server running on 18.0 or later in the absence of reverse DNS.

### Procedure

1. Download the NetWorker client and storage node package for Linux from the Online Support website to a temporary location.
2. Stop the NetWorker process using the `nsr_shutdown` command.
3. Install the client and storage node by running the `dpkg` command: `dpkg -i package package..`

```
dpkg -i lgtoc1nt_18.0_amd64.deb lgtonode_18.0_amd64.deb
```

4. Start the NetWorker daemons by typing the following command:

Initialization system	Command
<b>sysvinit</b>	<code>/etc/init.d/networker start</code>
<b>systemd</b>	<code>systemctl start networker</code>

5. Add the CloudBoost appliance using the NetWorker character-based interface (`nsradmin`) and then, label and mount the appliance.

## Report information on cloud backup

Use cloud backup information to monitor backup costs and help optimize the cloud backups.

Cloud backup information can be obtained from the following sources:

- Cloud backup and recover reports in the NMC.
- The `mminfo` command  
Use the `mminfo -avot` command to get information on how much data is consumed in a cloud backup. The *NetWorker Command Reference Guide* and the UNIX man pages provide more information about how to use the `mminfo` command.



# CHAPTER 9

## Perform a CloudBoost Appliance Recovery

This section includes the following topic:

- [Recovering CloudBoost Appliance](#).....78

## Recovering CloudBoost Appliance

You can recover the CloudBoost appliance by using the recovery metadata and private key.

### Before you begin

1. You must have the recovery metadata and the private key.
2. The recovery target appliance must be running the same version of the CloudBoost software as that of the failed appliance

### Procedure

1. Deploy a second CloudBoost appliance to restore the metadata from backups that are stored in the cloud.
2. Connect to CloudBoost through CLI and validate whether the date and time are in sync between the ESX server, CloudBoost virtual machine and the object store.
3. Use a web browser to sign in to the On-Prem CloudBoost Management Portal.
4. In the left menu, click **Appliance Recovery**.
5. Upload the metadata file and the private key and click **Preview**
6. Review the recovery information and click **Start Recovery**.

### After you finish

In the left menu, click CloudBoost Appliance and validate the configurations.

# CHAPTER 10

## Monitoring, Managing, and Supporting a CloudBoost Appliance

This section includes the following topics:

- [Monitoring CloudBoost](#) ..... 80
- [Upgrade a CloudBoost appliance](#) .....80
- [CloudBoost integration with EMC Secure Remote Services](#) ..... 81
- [Register CloudBoost with EMC Secure Remote Services](#)..... 82
- [Increase the CloudBoost appliance site cache](#).....83
- [Configuring average chunk size](#).....83
- [Specifications for the chunk size setting](#).....84

## Monitoring CloudBoost

CloudBoost is integrated with EMC Secure Remote Services (ESRS), which can be enabled to monitor the health of the appliances. If CloudBoost appliances are not registered with ESRS, you must monitor health, collect and review logs, and when necessary, contact Support.

## Upgrade a CloudBoost appliance

You can upgrade the CloudBoost appliance software in the CLI. During the upgrade, the CloudBoost appliance is unavailable. The appliance will restart after the upgrade is complete.

To review the latest CloudBoost appliance supported features, refer to the following content available on the Support website at <https://support.emc.com>:

- The [CloudBoost Release Notes](#) contains information about new features and changes, fixed problems, known limitations, environment, and system requirements for the latest release.
- The *Online Software Compatibility Guide* at <http://compatibilityguide.emc.com:8080/CompGuideApp/> provides a complete list of supported products and versions.

### Procedure

1. Connect to CloudBoost using an SSH client.

For more information about connecting to CLI, see [Connect to the CloudBoost CLI](#).

2. Login to CloudBoost as an administrator and run the following command:

```
upgrade appliance <URL upgrade path>.
```

3. Review the warning message, and then type **Yes**.

TheCloudBoost upgrade process initiates.

4. Monitor the progress of the upgrade.

---

### Note

During an upgrade, some services might temporarily be down or the appliance might go red while it restarts. The services and appliance will return to a normal state after the restart completes.

---

### Results

After the upgrade is complete, type `status` to view CloudBoost version.

### After you finish

By default, the recovery files are saved as `recovery.key` and `recovery.meta`. To prevent accidental overwrite of existing recovery files, it is recommended that you save the recovery files with a different file name.



**Note**

You must download the recovery information immediately. If you fail to download the recovery file, then you will be unable to perform a disaster recovery on the upgraded CloudBoost appliance.

---

## CloudBoost integration with EMC Secure Remote Services

EMC Secure Remote Services (ESRS) is a virtual appliance that enables two-way remote communication to monitor system health and to proactively communicate alerts and issues to Customer Support. ESRS is included at no extra charge in the enhanced or premium warranty or maintenance agreement.

### Registering EMC Secure Remote Services

When the CloudBoost appliance is registered with the EMC Secure Remote Services gateway, the appliance continuously communicates with EMC Secure Remote Services, sending status information and reports on a predetermined schedule. Appliance alerts from EMC Secure Remote Services appear in the On-Prem CloudBoost Management Console. When necessary, CloudBoost Technical Support is notified of issues and can open an SSH session with the appliance to obtain additional logs and reports.

**Auditing**

You can audit remote support activity, including the date and time of remote sessions, the ticket number, and the technician who provided the support.

**Registering**

You can allow or deny this remote activity for any reason. When a technical support agent starts a connection through EMC Secure Remote Services, an email that requests access is sent to you. You can choose to grant or deny the request.

If you choose not to register CloudBoost appliances with EMC Secure Remote Services, you must manually monitor the appliances. If any issues arise, contact Support.

### Installing the EMC Secure Remote Services gateway

You can install the EMC Secure Remote Services gateway version 3.6.0 or later in a VM separate from the CloudBoost appliance. After the CloudBoost appliance is registered in the On-Prem CloudBoost Management Console, you can then also register the appliance with EMC Secure Remote Services.

For information about installing the EMC Secure Remote Services gateway, refer to the EMC Secure Remote Services Virtual Edition topics at the following sites.

- [https://support.emc.com/products/37716\\_EM-SECURE-REMOTE-SERVICES-VIRTUAL-EDITION](https://support.emc.com/products/37716_EM-SECURE-REMOTE-SERVICES-VIRTUAL-EDITION)
- [https://support.emc.com/products/37716\\_EM-SECURE-REMOTE-SERVICES-VIRTUAL-EDITION/Topics/pg58757/](https://support.emc.com/products/37716_EM-SECURE-REMOTE-SERVICES-VIRTUAL-EDITION/Topics/pg58757/)

---

**Note**

At the CloudBoost CLI, when you register the appliance with EMC Secure Remote Services, you need to provide the SID from the email that is sent from EMC Secure Remote Services Support, along with the IP address or the URL and serial number. For information about registering the CloudBoost appliance with EMC Secure Remote Services, see [Register CloudBoost with EMC Secure Remote Services](#) on page 82.

---

## Register CloudBoost with EMC Secure Remote Services

You can register a CloudBoost appliance with the EMC Secure Remote Services gateway to enable two-way remote communication with Customer Support. EMC Secure Remote Services gateway monitors system health and communicates alerts and issues to Customer Support.

**Before you begin**

The EMC Secure Remote Services gateway must be installed, and the CloudBoost appliance must be registered in the On-Prem CloudBoost Management Console. Remote access must be enabled for the appliance.

---

**Note**

If a firewall exists between the CloudBoost appliance and the EMC Secure Remote Services gateway server, certain ports such as port 9443, must be open.

---

**Procedure**

1. Find the EMC Secure Remote Services SID in the email from EMC Secure Remote Services Support.
2. Have the IP address or URL and the serial number of the installed Secure Remote Services gateway available.
3. Connect to the CloudBoost CLI.
4. Type the following command:

```
support esrs register esrs_gateway username password sid
gateway_sn
```

where:

- *esrs\_gateway* is either the IP address or the FQDN for the EMC Secure Remote Services gateway virtual machine.
- *username* and *password* are the credentials that you used to set up the EMC Secure Remote Services gateway.
- *sid* is the EMC Secure Remote Services serial number that EMC Secure Remote Services Support provided in an email.
- *gateway\_sn* is the serial number for the EMC Secure Remote Services gateway.

---

**Note**

If you see this message, `Approval Request Pending - Contact EMC Customer Support`, contact Customer Support and ask for the device registration in EMC Secure Remote Services to be manually approved. After Support approves the request, you can run the command in step 4 again. After a device is successfully registered, you can also use the `status` command to verify the connection. At the bottom of the window, you will see a list of the EMC Secure Remote Services Server details.

---

**Results**

The CloudBoost appliance is registered with EMC Secure Remote Services, and continuous support monitoring begins.

## Increase the CloudBoost appliance site cache

After the CloudBoost appliance is deployed, you can increase the cache size by adding additional virtual data disks in a vCenter virtual machine configuration. Restart the CloudBoost appliance after you add the disks.

Any new data disk should be equal to the initial site-cache data-disk size.

- If a new data disk is smaller than the initial data disk size, CloudBoost generates a warning, and the disk is not added to the system.
- If a new data disk is bigger than the initial data disk size, CloudBoost generates a warning. The disk is added to the system, but the excess space is not used.
- The supported number of caching disks is 1, 2, 4, 8, 16, or 32.
  - If the number of available caching disks is not equal to the number of supported disks, site cache consumes the maximum supported disk count less than or equal to the available disks. The remaining disks are not used. For example, if there are 19 available caching disks, the caching server uses 16, and the rest are not used.
  - If there are 45 caching disks, the caching server uses 32, and the rest are not used.

---

**Note**

Additional data disks must be thick provisioned.

---

## Configuring average chunk size

You can adjust the chunk size in AWS and Azure to reduce API calls and improve backup, restore, and clone performance.

**Procedure**

1. Connect to the CloudBoost CLI.
2. Authenticate with the administrator password.
3. Check the current average chunk size appliance by typing the following command: `avgchunksize show`
4. To change the average chunk size, type the following command: `avgchunksize change xxx`

Where *xxx* is the new chunk size setting.

## Specifications for the chunk size setting

You can adjust the chunk size in AWS and Azure to reduce API calls and improve backup, restore, and clone performance.

Consider the following specifications for the chunk size setting:

- The maximum chunk size is 4 MB.
- The maximum average chunk size is 1 MB.
- The minimum average chunk size is 256 KB.
- If the average chunk size is increased, object sizes increase which reduces the overall API calls.
- The compression rate of each chunk is data dependent. Changing the chunk size results in smaller objects from higher compression ratios.

This table lists the workload chunk size specifications.

**Table 15** Specifications for work load types

Work load type	Chunk size			
	File system		SQL	
	256 K	1024 K	256 K	1024 K
Average chunk size in KB with Rabin finger print	282.29	1092.68	222	1347
Average chunk size in KB with transferring blob	294.2	1193.36	72	432

This table lists the compressed SQL data specifications.

**Table 16** Specification with compressed SQL data

Work load type	Chunk size	
	256 K	1024 K
Average chunk size in KB with Rabin finger print	310	1240
Average chunk size in KB with transferring blob	172	686