



EMC[®] Avamar[®]

Version 7.2

Operational Best Practices

302-001-931

REV 02

Copyright © 2001-2015 EMC Corporation. All rights reserved. Published in USA.

Published August, 2015

EMC 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. EMC Corporation 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 EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>).

EMC Corporation
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.EMC.com

CONTENTS

Figures		7
Tables		9
Preface		11
Chapter 1	Overview	15
	Guide organization	16
	Most important operational best practices	16
Chapter 2	Designing Avamar to Maximize System Availability	19
	Avamar architecture	20
	Stripes	20
	Avamar data server functions	20
	RAID, RAIN, replication, and checkpoints	21
	Redundant Array of Independent Disks (RAID)	21
	Redundant Array of Independent Nodes (RAIN)	21
	Best practices for RAIN	22
	Replication	22
	Best practices for replication	22
	Checkpoints	23
	Best practice for checkpoints	24
	Backing up clients in remote offices	24
	Best practice for backing up clients in remote offices	25
Chapter 3	Managing Capacity	27
	Impact of storage capacity when deploying a new Avamar system	28
	Impact of storage capacity on system performance	28
	Definitions of Avamar server capacities	28
	Avamar capacity thresholds	28
	Impact of capacity on various operations	29
	Best practices for capacity management	29
	Proactive steps to manage capacity	30
	Steps to recover from capacity issues	31
	Steady state system	31
Chapter 4	Scheduling	33
	Avamar clients	34
	Restrictions and limitations	34
	Best practices for Avamar clients	35
	Scheduling daily activities	38
Chapter 5	Defining Domains, Groups, and Policies	39

	Best practices for making management policy decisions	40
	Defining domains	40
	Best practices for defining domains	40
	Defining groups	40
	Best practices for defining groups	41
	Defining datasets	41
	Best practice for defining datasets	41
	Defining schedules and retention policies	41
	Best practices for defining schedules	41
	Best practices for setting up retention policies	42
Chapter 6	Daily Monitoring of Backup Infrastructure and Operations	43
	Monitoring the Avamar system	44
	Monitoring the Avamar system backup operations	45
	Closely monitor daily backup activities	45
	Best practices for monitoring daily backup activities	46
	Closely monitor nightly replication	47
Chapter 7	Tuning Performance	49
	Demand-paging cache	50
	File cache	51
	Demand-paging cache example scenario	51
	Cache migration after upgrading the client	52
	Caching options	52
	Client caching	53
	Cache information in the avtar logs	54
	Using cacheprefix	54
	Example for using the cacheprefix attribute	55
	Configuring custom hash settings for Microsoft databases	56
	Tuning replicator	56
	Backup and restore performance	56
Chapter 8	Understanding DPN Summary Reports	57
	DPN Summary report	58
	Running a DPN Summary report	58
	Example DPN Summary entry	59
	Background on backups	62
	Dataset size	62
	Modified files	62
	Summary of key DPN summary terms	64
	Definition of commonality	65
	Avamar backup compared to incremental tape backup	65
	Definition of terms during restore activities	65
Chapter 9	Protecting Avamar Desktop/Laptop Clients	67
	About Avamar Desktop/Laptop	68
	Deploy additional Avamar servers for Desktop/Laptop clients	68
	Best practice for deploying Desktop/Laptop clients	68
	Create a dataset to back up only user data	68
	Best practices for creating a dataset	68
	Exclusions for Windows computers	69

	Exclusions for Mac computers.....	70
	Minimize the number of exclude and include lists.....	70
	Dataset caveat.....	70
	Keep the initial client backups to a manageable number.....	71
	Best practices for first-time backups.....	71
	Strategy for performing first-time backups.....	71
	Strategy for setting up backup groups.....	72
	Activating clients by using the Avamar Client Manager.....	72
	Consider node size when configuring Avamar servers.....	73
	Determine the backup window.....	73
	Best practice for determining the backup window.....	74
	Schedule the backup window.....	74
	Best practice for scheduling backups.....	74
	Adjust runtime for daily maintenance tasks.....	74
	Best practice for scheduling daily maintenance tasks.....	75
	Do not run client utilities during the backup window.....	75
	Run backups more frequently than the retention policy.....	75
	Prevent backups on a wireless connection.....	75
	Clearing the Back Up On Wireless option on Windows.....	76
	Clearing the Back Up On Wireless option on Mac.....	76
	Manage storage capacity for Desktop/Laptop clients.....	76
	Ensure adequate initialization time for Wake-on-Lan backups.....	76
	Best practices for Wake-on-Lan backups.....	77
Chapter 10	Other Avamar Administration Best Practices	79
	Protecting the Avamar server.....	80
	Use of an uninterruptible power supply with Avamar.....	80
	Best practices for using a UPS with Avamar.....	80
	Changing passwords.....	80
	Best practice for default passwords.....	81
	Using Avamar Client Manager.....	81
	Best practice for using Avamar Client Manager.....	82
	Enabling the Email Home feature.....	82
	Using EMC Secure Remote Support solution.....	82
	Best practice for EMC Secure Remote Support.....	82
	Assigning users.....	82
	Best practices for assigning users.....	83
Chapter 11	Using Data Domain Systems	85
	Network bandwidth recommendations	86
	Use the iperf utility to test the network bandwidth.....	86
	Recommended network bandwidth.....	87
	Example iperf utility sessions.....	87
	Configuration best practices.....	89
	Use fully qualified domain names.....	89
	Review the amount of files in the MTree.....	89
	Do not modify the MTree.....	90
	Specify the maximum number of data streams.....	90
	Evaluate storage requirements.....	90
	Synchronize the time.....	90
	Restore backups with the Use SQL REPLACE option.....	90
	Space requirements for replication configurations.....	90
	Data movement policies.....	91

Chapter 12	Using Isilon Storage Devices	93
	Isilon overview	94
	Configuration	94
	Guidelines for configuring the Avamar environment	95
	Backups	96
	Best practices for backing up Isilon	97
	Exclude lists for Isilon backups	97
	Multi-streaming backups	100
	Replication target backups	102
	Backups after renaming a TLD	102
	Backups after moving a TLD inside an existing TLD	102
	Restores	102
	Best practice for restoring Isilon	102
	Multi-streaming restores	102
	Cross-platform restores	102
	Concurrent backups and restores	103
	Performance scalability metrics	103
	Isilon backups to and restores from Avamar Data Store	103
	Summary of performance testing	104

FIGURES

1	Backup/Maintenance windows	38
2	Edit Advanced Retention Policy dialog box.....	42
3	Monolithic cache compared to page cache.....	50
4	Data flow to server.....	63
5	Avamar commonality diagram.....	64
6	Isilon configuration with Avamar and Data Domain.....	95
7	Isilon root directory and three TLDs.....	96
8	Backup selection comprising TLDs and subdirectories.....	101

TABLES

1	Revision history	11
2	Best practices guide's organization.....	16
3	Lifecycle phases.....	16
4	Types of stripes.....	20
5	Avamar server operational functions.....	21
6	Remote office backups	24
7	Capacity thresholds.....	29
8	Known restrictions and limitations for planning and designing the Avamar system.....	34
9	Ways to monitor the Avamar system	44
10	Client messages for client backups	45
11	Avamar Administrator reports for client backups.....	46
12	Options for avtar to control demand-paging cache.....	52
13	Segregating data into separate datasets.....	55
14	DPN Summary column descriptions.....	59
15	Desktop/Laptop file types to include in a dataset.....	69
16	Desktop/Laptop file types to exclude from a dataset	69
17	Avamar user accounts and SSH keys.....	81
18	Requirements for Isilon NAS device and Avamar configuration.....	95
19	Configuration for testing Isilon backups to and restores from an Avamar Data Store....	103
20	Performance results for level 0 backups to an Avamar Data Store.....	104
21	Performance results for level 1 backups to an Avamar Data Store.....	104
22	Performance results for restore operations from an Avamar Data Store.....	104

PREFACE

As part of an effort to improve its product lines, EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

Contact your EMC technical support professional if a product does not function properly or does not function as described in this document.

Note

This document was accurate at publication time. Go to EMC Online Support (<https://support.EMC.com>) to ensure that you are using the latest version of this document.

Purpose

This guide describes operational best practices for both single-node and multi-node servers in small and large heterogeneous client environments. This guide does not provide introductory materials for basic Avamar technology or delivery methods.

Audience

The intended audience of this document is experienced UNIX, Linux, and Windows system administrators who will deploy and operate Avamar servers.

Revision history

The following table presents the revision history of this document.

Table 1 Revision history

Revision	Date	Description
01	June, 2015	Initial release of Avamar 7.2.
02	August 2015	GA release of Avamar 7.2. Minor updates.

Related documentation

The following EMC publications provide additional information:

- *EMC Avamar Administration Guide*
- *EMC Avamar Backup Clients User Guide*
- *EMC Avamar Data Store Customer Service Guide*
- *EMC Avamar Data Store Single Node Customer Installation Guide*
- *EMC Avamar for Exchange VSS User Guide*
- *EMC Avamar for IBM DB2 User Guide*
- *EMC Avamar for Lotus Domino User Guide*
- *EMC Avamar for SharePoint VSS User Guide*
- *EMC Avamar for Oracle User Guide*
- *EMC Avamar for SAP with Oracle User Guide*
- *EMC Avamar for SQL Server User Guide*

- *EMC Avamar for Sybase ASE User Guide*
- *EMC Avamar for VMware User Guide*
- *EMC Avamar for Windows Server User Guide*
- *EMC Avamar Management Console Command Line Interface (MCCLI) Programmer Guide*
- *EMC Avamar NDMP Accelerator for EMC NAS Systems User Guide*
- *EMC Avamar Product Security Guide*
- *EMC Avamar Release Notes*
- *EMC Backup & Recovery Manager User Guide*
- White paper: *Efficient Data Protection with EMC Avamar Global Deduplication Software - Technology Concepts and Business Considerations*
- White paper: *Optimized Backup and Recovery for VMware® Infrastructure with EMC Avamar*

Special notice conventions used in this document

EMC uses the following conventions for special notices:

NOTICE

Addresses practices not related to personal injury.

Note

Presents information that is important, but not hazard-related.

Typographical conventions

EMC uses the following type style conventions in this document:

Bold	Use for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)
<i>Italic</i>	Use for full titles of publications referenced in text
Monospace	Use for: <ul style="list-style-type: none"> • System code • System output, such as an error message or script • Pathnames, file names, prompts, and syntax • Commands and options
<i>Monospace Italic</i>	Use for variables
Monospace bold	Use for user input
[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Where to get help

The Avamar support page 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 EMC Customer Support.

To access the Avamar support page:

1. Go to <https://support.EMC.com/products>.
2. Type a product name in the **Find a Product** box.
3. Select the product from the list that appears.
4. Click the arrow next to the **Find a Product** box.
5. (Optional) Add the product to the **My Products** list by clicking **Add to my products** in the top right corner of the **Support by Product** page.

Documentation

The Avamar product documentation provides a comprehensive set of feature overview, operational task, and technical reference information. Review the following documents in addition to product administration and user guides:

- Release notes provide an overview of new features and known limitations for a release.
- Technical notes provide technical details about specific product features, including step-by-step tasks, where necessary.
- White papers provide an in-depth technical perspective of a product or products as applied to critical business issues or requirements.

Knowledgebase

The EMC Knowledgebase contains applicable solutions that you can search for either by solution number (for example, esgxxxxx) or by keyword.

To search the EMC Knowledgebase:

1. Click the **Search** link at the top of the page.
2. Type either the solution number or keywords in the search box.
3. (Optional) Limit the search to specific products by typing a product name in the **Scope by product** box and then selecting the product from the list that appears.
4. Select **Knowledgebase** from the **Scope by resource** list.
5. (Optional) Specify advanced options by clicking **Advanced options** and specifying values in the available fields.
6. Click the search button.

Online communities

Visit EMC Community Network at <http://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 EMC products.

Live chat

To engage EMC Customer Support by using live interactive chat, click **Join Live Chat** on the **Service Center** panel of the Avamar support page.

Service Requests

For in-depth help from EMC Customer Support, submit a service request by clicking **Create Service Requests** on the **Service Center** panel of the Avamar support page.

Note

To open a service request, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

To review an open service request, click the **Service Center** link on the **Service Center** panel, and then click **View and manage service requests**.

Facilitating support

EMC recommends that you enable ConnectEMC and Email Home on all Avamar systems:

- ConnectEMC automatically generates service requests for high priority events.
- Email Home emails configuration, capacity, and general system information to EMC Customer Support.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to DPAD.Doc.Feedback@emc.com.

Please include the following information:

- Product name and version
- Document name, part number, and revision (for example, 01)
- Page numbers
- Other details that will help us address the documentation issue

CHAPTER 1

Overview

This chapter includes the following topics:

- [Guide organization](#) 16
- [Most important operational best practices](#) 16

Guide organization

The following table shows the best practices guide's organization.

Table 2 Best practices guide's organization

Subject matter	Chapters
Core EMC® Avamar® system functions	<ul style="list-style-type: none"> • Designing Avamar to Maximize System Availability on page 19 • Managing Capacity on page 27 • Scheduling on page 33 • Defining Domains, Groups, and Policies on page 39 • Daily Monitoring of Backup Infrastructure and Operations on page 43
Tuning the Avamar system	<ul style="list-style-type: none"> • Tuning Performance on page 49 • Understanding DPN Summary Reports on page 57
Avamar Desktop/Laptop clients	<ul style="list-style-type: none"> • Protecting Avamar Desktop/Laptop Clients on page 67
Other Avamar administration functions	<ul style="list-style-type: none"> • Other Avamar Administration Best Practices on page 79
Data Domain systems	<ul style="list-style-type: none"> • Using Data Domain Systems on page 85
Isilon systems	<ul style="list-style-type: none"> • Using Isilon Storage Devices on page 93

The following table describes Avamar server lifecycle phases that the contents in this guide covers.

Table 3 Lifecycle phases

Lifecycle phase	Description
Planning and design	Topology and architecture options, risks and limitations, and any other planning and design issues that you must consider before implementing the design.
Implementation	Installation options and directions for testing Avamar components after the installation is complete.
Daily operations	Normal management of Avamar server capacity, performance optimization of backups and replication, and daily monitoring of the Avamar infrastructure and operations.

Most important operational best practices

Here are the most important best practices to understand and follow:

- Check the EMC Online Support (<https://support.EMC.com>) for the most current version of the *EMC Avamar Operational Best Practices*.
- Set up a correct daily operational schedule for the Avamar server.
- Proactively assess and correct systematic issues.
- Deploy a uninterruptible power supply (UPS) for the Avamar server hardware to protect against data loss caused by unplanned power outages.
- Do not add clients to the root (/) domain in the Avamar server.
- Understand how to monitor and manage the storage capacity of the Avamar server on a daily basis.
- Minimize the number of groups used to back up clients. Schedule backups during the server's backup window so that the backups do not overlap with daily maintenance jobs.
- Monitor the Avamar server on a daily basis. Interpret all system warnings and errors.
- Investigate all failed backups, missing clients, and backups that completed with exceptions.
- Protect the Avamar server from the Internet by providing full firewall protection.
- Change all factory default passwords including the default SSH keys. Do not change the passwords for the backuponly, restoreonly, and backuprestore software application users.
- Ensure every administrator logs in to the Avamar server with a unique username.
- Enable the Email Home capability.
- Check the network bandwidth before adding a Data Domain system to an Avamar configuration.
- Do not use an Avamar single node server for any other purpose other than for storing client backups.

Modifying the Avamar system in any way other than as instructed by product documentation or authorized EMC personnel is strictly forbidden. This includes installing third-party software on the node; creating, modifying or deleting any file or directory in an Avamar system; or changing any configuration settings in the hardware, firmware, or operating system.

The chapters that follow provide more details about these best practices and other best practices.

CHAPTER 2

Designing Avamar to Maximize System Availability

This chapter includes the following topics:

- [Avamar architecture](#)..... 20
- [RAID, RAIN, replication, and checkpoints](#)..... 21
- [Backing up clients in remote offices](#)..... 24

Avamar architecture

To ensure the long-term reliability, availability, and supportability of the Avamar server, you must design it carefully.

Several processes run on the Avamar server nodes. Key processes include:

- Avamar Administrator server on the utility node.
- Avamar data server on all active storage nodes.

GSAN (Global Storage Area Network) is another term that refers to the Avamar data server.

The Avamar data server stores, processes, and manages the variable-sized chunks that the client sends during a backup. An average size chunk is about 10 KB depending on the customer data. Through the patented deduplication technology, backups send only unique data chunks to the Avamar data server.

Stripes

The term “stripe” refers to the container an Avamar data server uses to manage the data in the system. Stripes are files of various sizes that are based on the kind of stripe.

Each stripe has a unique name. The Avamar server can identify and access a stripe by name only. The following table describes four types of stripes.

Table 4 Types of stripes

Stripe	Description
Atomic data	Contains data that originates on the customer system and is read during a backup.
Composite	Contains references to other composite or atomic stripes, and provides the means to build trees that can arbitrarily represent large amounts of data. References are SHA-1 hashes.
Index	Maps a hash to the stripe that contains corresponding data. The index is the essence of a “content addressed” store.
Parity	Provides simple XOR parity that can be used to reconstruct data when a failure occurs. Every stripe belongs to a parity group that protects it when you use RAIN. A protected stripe is called a “safe” stripe.

Avamar data server functions

The Avamar data server is a high-transaction-rate database-like application that is optimized to store and manage billions of variable-sized objects in parallel across all active storage nodes.

The Avamar server performs several functions throughout each day. The following table describes the major operational functions.

Table 5 Avamar server operational functions

Function	Description
Backup	Supports the backup operation by receiving, processing, and storing the backup data that Avamar clients send. During a backup, the Avamar server interacts with the client to ensure that the client sends only unique data chunks to the server.
Restore	Restores the data stored on the Avamar server to the Avamar client.
Checkpoint	Creates consistent point-in-time images (checkpoints) every day. Avamar uses checkpoints as rollback points to recover from various issues, such as sudden power loss.
hfscheck	Validates one of the checkpoints every day through a process called <code>hfscheck</code> .
Garbage collection	Deletes the chunks of data that are no longer referenced within the backups stored on the system.
Replication	Supports daily replication of the backups.
Precrunching	Prepares stripes throughout the day that Avamar reuses during backup operations. During the precrunching process, the server selects the emptiest stripes, those that contain more empty space than the data partitions (by percentage), and defragments them. The precrunching process leaves contiguous space for new data.

The Avamar server requires adequate CPU, memory, and I/O resources to perform these functions throughout the day. Avamar performs extensive qualification testing of all approved platforms to ensure that the resources available are adequate to meet long-term reliability, availability, and supportability requirements.

RAID, RAIN, replication, and checkpoints

The Avamar system provides up to four levels of systematic fault tolerance: RAID, RAIN, replication, and checkpoints.

Redundant Array of Independent Disks (RAID)

All standard Avamar server node configurations use RAID to protect the system from disk failures. RAID enables you to hot swap the hard disk drives that have been the highest failure rate hardware items in Avamar servers.

Failed drives impact I/O performance and affect the Avamar server performance and reliability. In addition, RAID rebuilds can significantly reduce the I/O performance, which can adversely impact the performance and reliability of the Avamar server.

Redundant Array of Independent Nodes (RAIN)

RAIN provides the means for the Avamar server to continue to operate even when a node fails. If a node fails, you use RAIN to reconstruct the data on a replacement node.

In addition to providing failsafe redundancy, you can use RAIN to rebalance the capacity across the nodes after you expand the Avamar server by adding new nodes. The ability to manage the capacity of the system as the amount of data added to the system continues

to increase is a critical feature. Except for two-node systems, RAIN protection is enabled in multi-node Avamar servers. Single-node servers do not use RAIN.

Best practices for RAIN

- Always enable RAIN for all configurations other than single-node servers. Minimum RAIN configuration is a 1x3 system (three active storage nodes plus a utility node and optionally, a spare node).

Note

Spare nodes are optional in ADS Gen4 systems that run Avamar 7.2.

Double-disk failures on a node or a complete RAID controller failure can occur. Either of these failures can corrupt the data on a node. Without RAIN, the only recourse is to reinitialize the entire system and replicate the data back from the replication target.

Note

ADS Gen4 with Avamar 7.2 does not support new installs of 1x2 systems.

- When deploying single-node servers, you must replicate the data on the servers to ensure that you protect the data.

Note

The business edition of Gen4 RAID6 nodes do not require replication.

Non-RAIN servers have no data redundancy and any loss of data requires that the system be reinitialized.

- Limit initial configurations to 12 to 14 active storage nodes so that you can add nodes later if needed to recover from high-capacity utilization situations.

Replication

The Avamar system can efficiently replicate data from one Avamar server to another on a scheduled basis. The replication server ensures complete data recovery if you lose the primary backup Avamar server.

Replication is useful for more than recovering a single client. Replication moves data to another system that can be used for data recovery if an unexpected incident occurs. Replication is, by far, the most reliable form of redundancy that the system can offer because it creates a logical copy of the data from the replication source to the destination. Replication does not create a physical copy of the blocks of data. Any corruptions, whether due to hardware or software, are less likely to propagate from one Avamar server to another. In addition, multiple checks of the data occur during replication to ensure that only uncorrupted data is replicated to the replication target.

If maximizing the availability of the backup server for backups and restores is important, set up a replication system as quickly as possible.

Best practices for replication

- Protect the data on the Avamar server by replicating the data to another Avamar server.

- Use default standard replication, also known as “root-to-REPLICATE” replication, to do the following:
 - Provide the flexibility to configure replicated Avamar servers in a wide variety of ways
 - Have full visibility into all the backups that you replicate from one Avamar server to another

Standard replication also supports the ability to replicate the contents of many replication source Avamar servers to a single large replication destination (many-to-one), or to cross-replicate the contents of a couple of Avamar servers to each other. At any time, you can browse the contents of the `/REPLICATE` domain on the replication destination and see all the backups that have been replicated for each account.

- Ensure that available network bandwidth is adequate to replicate all the daily changed data within a four-hour window so that the system can accommodate peaks of up to eight hours per day. The replicator can use 60% to 80% of the total available bandwidth when WAN bandwidth is the performance bottleneck. The *EMC Avamar Administration Guide* contains more information about setting up replication to best use the system bandwidth.
- When defining daily replication, avoid using the `--include` option. Use the `--include` option to perform only selective replication under certain conditions. The use of the `--include` option to list clients for replication is prone to error. To add a new client to the Avamar server, edit the `repl_cron.cfg` file and add a new `--include` option for the client. After adding the client to the `repl_cron.cfg` file, the client data is replicated.
- Use the `--exclude` option only if you decide to selectively exclude a high change-rate or low-priority client from the nightly replication.
- When configuring replication, always set the `--retention-type` option to replicate all retention types (none, daily, weekly, monthly, and yearly). If you leave out retention type “none” from the replication, the hourly Avamar Administrator server backups are not replicated. To perform a full disaster recovery of the replication source Avamar server requires these system backups.

Checkpoints

Checkpoints provide redundancy across time. Checkpoints enable you to recover from operational issues. For example:

- Trying to back up a client that is too large to fit in the available remaining capacity.
- Accidentally deleting a client and all the associated backups.

In addition, checkpoints enable you to recover from certain kinds of corruption by rolling back to the last validated checkpoint.

Checkpoints are an effective way to revert the system back to an earlier point in time. Checkpoints, like all other forms of redundancy, require disk space. The more checkpoints you retain, the larger the checkpoint overhead.

Best practice for checkpoints

Leave the checkpoint retention policy at the default values. The default is set to retain the last two checkpoints, whenever created, and the last validated checkpoint.

Note

During certain support actions, EMC Customer Service might temporarily change the checkpoint retention policy to ensure that certain critical checkpoints are retained during the support action. After EMC Customer Service completes the support action, restore the checkpoint retention policy to the default setting.

Backing up clients in remote offices

When you back up clients in a remote office, consider the following options:

- Option 1—Remote office with a small Avamar server and centralized replication destination

Is it better to back up clients in a remote office to a small Avamar server in the remote office, and then replicate data to a large centralized Avamar server?

- Option 2—Large centralized Avamar server and large centralized replication destination

Is it better to back up clients in a remote office directly to a large centralized Avamar server, and then replicate data to another large centralized Avamar server?

The following table lists factors to help you determine a remote backup strategy.

Table 6 Remote office backups

Factor	Scenario
Recovery time objective (RTO)	To restore data, the Avamar server compresses the restored data and sends it to the Avamar client where the data is then uncompressed. The restore process does not perform deduplication on the restored data. The primary advantage of backing up data to a remote Avamar backup server (Option 1) is that you can restore the data directly from the server across the local area network to the client. This advantage is important if you must satisfy an RTO requirement.
Server administration	<p>The amount of administration and support you require in a remote office is proportional to the number of Avamar servers you deploy in an environment. For example, 10 single-node servers you deploy as remote Avamar backup servers require considerably more administration and support than a single 1x8+1 multi-node configuration of 10 nodes that functions as a centralized Avamar backup server.</p> <hr/> <p>Note</p> <p>The single 1x8+1 multi-node configuration comprises eight active storage nodes, one utility node, and one spare.</p>
IT resources	If you deploy a remote Avamar backup server at a remote office, adequate IT resources for disaster recovery restores might not be available. In this case, Option 2 might be appropriate. A centralized IT staff can perform disaster

Table 6 Remote office backups (continued)

Factor	Scenario
	recovery restores to replacement hardware at the central site, and then ship the fully-configured replacement client to the remote site.
Exchange Server	If you have a Microsoft Exchange Server in the remote office with bandwidth limitations, Option 1 might be more practical. Back up the Exchange Server's storage group or database to a local Avamar server.
Large multi-node servers	If you require large multi-node servers to back up all data in a remote office, the cost of deploying, managing, and supporting remote Avamar servers might be the same as using centralized Avamar backup servers.

If the environment's WAN throughput is a bottleneck, the time a nightly replication requires in Option 1 is approximately the same as the time a backup requires in Option 2. The trade-off then becomes RTO compared to the additional cost of deploying, managing, and supporting multiple Avamar server instances.

Best practice for backing up clients in remote offices

Unless you cannot meet the RTO, design the system so that clients first back up directly to a large, active, and centralized Avamar server. Then replicate the data to another large centralized Avamar server.

CHAPTER 3

Managing Capacity

This chapter includes the following topics:

- [Impact of storage capacity when deploying a new Avamar system](#) 28
- [Impact of storage capacity on system performance](#) 28

Impact of storage capacity when deploying a new Avamar system

When you deploy a new Avamar system, the server usually fills rapidly for the first few weeks because nearly every client that you back up contains unique data. You can best leverage the Avamar commonality feature when you back up other similar clients or back up the same clients at least once.

After the first backup of a client, the Avamar system backs up less unique data during subsequent backups. When initial backups are complete and the maximum retention periods are exceeded, the system achieves steady state capacity utilization. Steady state capacity utilization is when the amount of new data sent to the Avamar server each day is about the same amount as what is freed during the maintenance windows.

Successfully achieving steady state capacity utilization is important for fixed-capacity systems such as single-node servers.

Impact of storage capacity on system performance

When managing an Avamar server, you can improve the long-term reliability, availability, and manageability of the Avamar server if you do either of the following:

- Minimize the average daily data change rate of the Avamar backup clients. The *EMC Avamar Administration Guide* provides more information.
- Reduce the per-node capacity within the Avamar server:
 - Reduce backup retention
 - Ensure daily maintenance jobs run regularly
 - Add more nodes to the Avamar server

Note

Many of the best practices in this guide can help you understand the average daily change rate and help you manage the per-node capacity.

Definitions of Avamar server capacities

Storage subsystem (GSAN) capacity is the total amount of commonality factored data and RAIN parity data (net after garbage collection) on each data partition of the server node. The GSAN process measures and reports this amount.

The administrator of the Avamar server can control this reported capacity:

- First, by changing the dataset definitions, retention policies, or even the clients that the server backs up.
- Next, by ensuring that a garbage collection operation runs regularly to remove expired data.

Operating system capacity is the total amount of data in each data partition, as measured by the operating system. This amount is not particularly useful to an external observer because the server manages disk space itself.

Avamar capacity thresholds

The GSAN changes behavior as the various capacities increase. The following table describes the behavior of key capacity thresholds.

Table 7 Capacity thresholds

Threshold	Default values	Capacity used for comparison	Behavior
Capacity warning	80% of read-only threshold	GSAN	The Management Console Server issues a warning event when the GSAN capacity exceeds 80% of the read-only limit.
Healthcheck limit	95% of read-only threshold	GSAN	If the GSAN capacity reaches the healthcheck limit, the Avamar server allows existing backups to complete, but suspends all new backup activity. Avamar sends a notification in the form of a pop-up alert when you log in to Avamar Administrator. You must acknowledge this alert before the system can resume activity.
Server read-only limit	100% of read-only threshold, which is set to a prespecified percentage of available hard drive capacity	GSAN	If the GSAN capacity on any data partition on any node exceeds the read-only threshold, the Avamar server transitions to read-only state to prevent the addition of new data. You can view the server utilization value on the Server Management tab (Avamar Administrator > Server > Server Management). The reported value represents the average utilization relative to the read-only threshold.
System too full to perform garbage collection	85% of available hard drive capacity	Internal GSAN calculation	If the GSAN determines that the space available on any data partition on any node exceeds the disknpgc configuration threshold, a garbage collection operation does not run. The operation fails with the error message <code>MSG_ERR_DISKFULL</code> .

Impact of capacity on various operations

As the amount of data stored in the Avamar server increases, the maintenance operations take longer to complete. Most notable is the garbage collection activity.

Any variations with incoming data or with daily maintenance routines can cause the system to become read-only or additional maintenance routines to fail.

Best practices for capacity management

- Monitor and manage the storage capacity of the Avamar server on a daily basis.
- Limit storage capacity usage to 80% of the available GSAN capacity.
- Monitor all variations with incoming data to prevent the system from becoming read-only.

Note

A system becomes read-only when storage capacity exceeds 80%. At this point the system capacity is considered full.

- Monitor all variations with maintenance jobs to prevent these jobs from failing.

Proactive steps to manage capacity

You receive a warning when the server storage capacity exceeds 80% of the read-only threshold.

When the capacity exceeds 80%, perform the following steps.

Procedure

1. Stop adding new clients to the system.
2. Reassess retention policies to see if you can decrease the retention, and therefore, reduce the capacity use.
3. Investigate the possibility that backups are preventing a garbage collection operation from starting:

- a. Use `dumpmaintlogs --types=gc` to view logs for the garbage collection operation.

Look for either of the following error messages in the garbage collection log:

```
MSG_ERR_BACKUPSINPROGRESS
```

```
garbage collection skipped because backups in progress
```

- b. Use `capacity.sh` to:

- Assess the data change rate in the environment.
- Assess the garbage collection effectiveness.
- Ensure that the system is running in steady state.
- Identify the three highest change rate clients.

The `capacity.sh` command displays output similar to the following:

```
admin@avamar-1:~/>:capacity.sh
```

Date	New Data	#BU	Removed	#GC	Net Change
2014-11-06	4888 mb	6	-1 mb	4	4887 mb
2014-11-07	1232 mb	9	0 mb		1232 mb
2014-11-08	63902 mb	9	-2 mb	4	63900 mb
2014-11-12	1158 mb	4	0 mb		1158 mb
2014-11-13	497 mb	7	-1 mb	1	496 mb
2014-11-14	1661 mb	8	-1 mb	1	1660 mb
2014-11-15	4772 mb	10	-1 mb	1	4771 mb
2014-11-16	781 mb	8	-268 mb	1	513 mb
2014-11-17	701 mb	9	0 mb	1	701 mb
2014-11-18	369 mb	7	0 mb	1	369 mb
2014-11-19	503 mb	9	0 mb	1	503 mb
2014-11-20	1630 mb	7	0 mb		1630 mb
Average	6841 mb		-22 mb		6818 mb
Top 5 Capacity Clients			Added	% of Total	ChgRate
client1			68405 mb	83.3%	3.022%
client2			4571 mb	5.6%	1.851%
client3			3844 mb	4.7%	2.592%
client4			3062 mb	3.7%	1.914%
client5			1738 mb	2.1%	0.128%
Total for all clients			82100 mb	100.0%	0.016%

where:

- The New Data column lists the daily amount of data added to the Avamar server.
- The BU# column lists the number of backup or replication jobs that occur daily.
- The Removed column lists the amount of data that GC recovers.
- The GC column lists the number of times GC runs each day.
- The Net Change column lists the amount of data added or removed each day.

The final summary lists the clients with the highest change rates.

Steps to recover from capacity issues

After the Avamar server capacity exceeds the warning threshold and approaches the `diskreadonly` limit, perform one or more of the following actions.

Procedure

- Follow the steps described in [Proactive steps to manage capacity on page 30](#).
- Acknowledge healthcheck limit events. When the Avamar server reaches the healthcheck limit, Avamar suspends all new backup activity until you acknowledge the healthcheck limit event.
- Contact EMC Customer Service if the Avamar server transitions to a read-only state.
- Replicate the data to another server temporarily, and then replicate the data back after you reinitialize the server. Replication creates a logical copy of the data, which compacts all the data onto fewer stripes.
- Add nodes and rebalance the capacity if the Avamar server is a multi-node server that uses RAIN. If the server has eight or more active storage nodes, to noticeably reduce the capacity per node, add two or more nodes at a time rather than adding only one node. You can add a maximum of four nodes at a time.

Steady state system

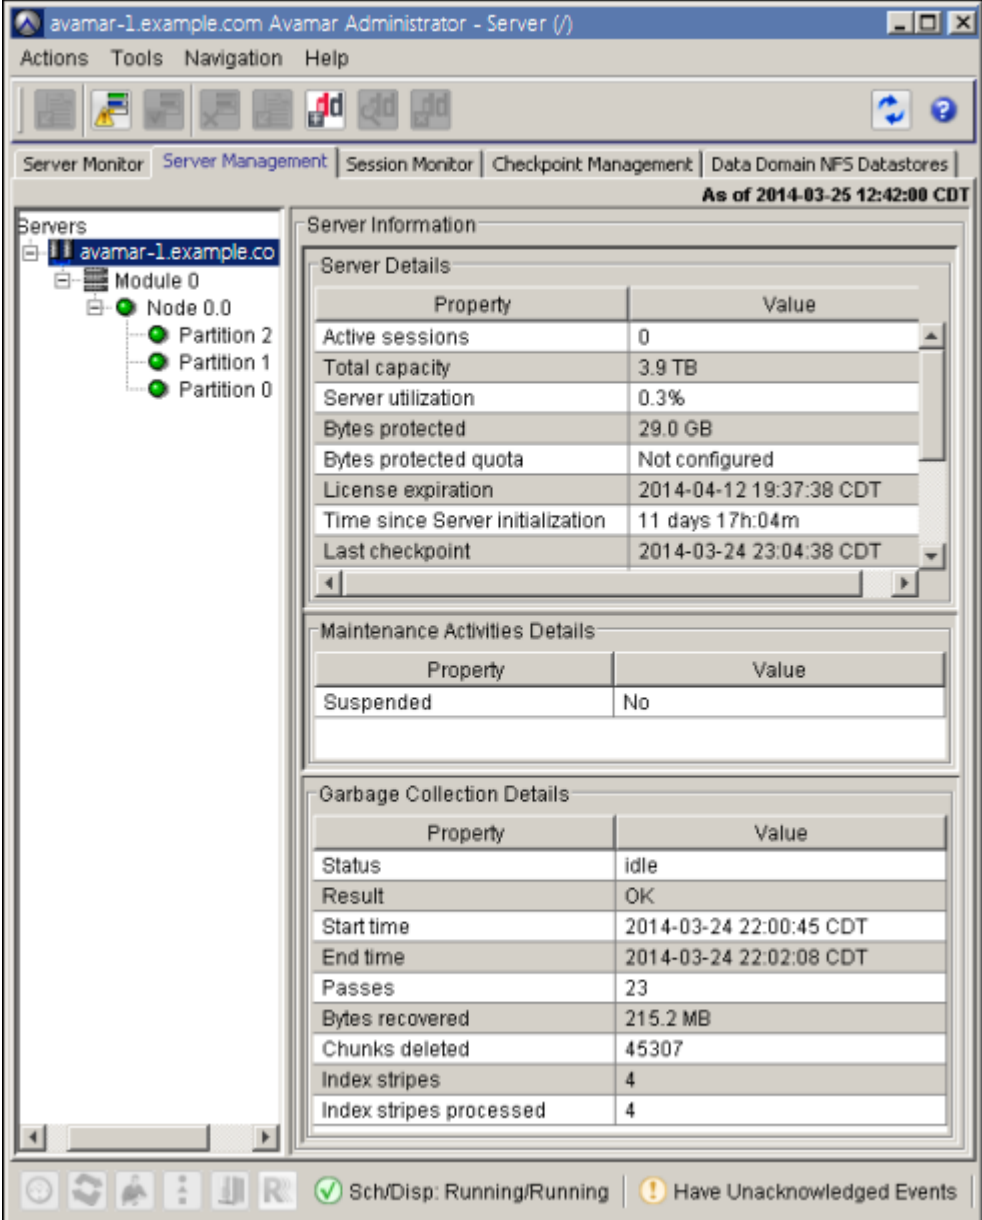
Typically, an Avamar system achieves steady state shortly after the longest retention period for the backups. For example, if you retain all daily backups for 30 days and all monthly backups for 3 months, the system starts to operate in steady state about 3½ to 4 months after you add the last client to the system. A slight delay occurs before achieving steady state because the garbage collection process requires several passes before it reaches the bottom of the file system tree. Garbage collection finds orphaned chunks in the upper levels first before removing orphaned data in the lower levels of the file system.

After the system achieves steady state, verify server utilization.

Procedure

1. Schedule activities so that all backups and maintenance tasks run successfully.
2. In Avamar Administrator, click the **Server** launcher button.
The **Server** window appears.
3. Click the **Server Management** tab.

Avamar server information appears in the window.



avamar-1.example.co Avamar Administrator - Server (/)

Actions Tools Navigation Help

Server Monitor Server Management Session Monitor Checkpoint Management Data Domain NFS Datastores

As of 2014-03-25 12:42:00 CDT

Servers

- avamar-1.example.co
 - Module 0
 - Node 0.0
 - Partition 2
 - Partition 1
 - Partition 0

Server Information

Server Details

Property	Value
Active sessions	0
Total capacity	3.9 TB
Server utilization	0.3%
Bytes protected	29.0 GB
Bytes protected quota	Not configured
License expiration	2014-04-12 19:37:38 CDT
Time since Server initialization	11 days 17h:04m
Last checkpoint	2014-03-24 23:04:38 CDT

Maintenance Activities Details

Property	Value
Suspended	No

Garbage Collection Details

Property	Value
Status	idle
Result	OK
Start time	2014-03-24 22:00:45 CDT
End time	2014-03-24 22:02:08 CDT
Passes	23
Bytes recovered	215.2 MB
Chunks deleted	45307
Index stripes	4
Index stripes processed	4

Sch/Disp: Running/Running | Have Unacknowledged Events

4. Verify that **Server utilization** in the **Server Details** table is at or below 80%.

CHAPTER 4

Scheduling

This chapter includes the following topics:

- [Avamar clients](#)..... 34
- [Scheduling daily activities](#).....38

Avamar clients

Avamar client agents are applications that run natively on the client systems. The Avamar client software is composed of at least two executable programs: `avagent` and `avtar`. The `avagent` program runs as a service on the client and establishes and maintains communication with the Avamar Administrator server.

When the Avamar Administrator server queues a backup work order, the Avamar Administrator server pages `avagent`. If the Avamar Administrator server cannot establish a connection with the client, `avagent` polls the Avamar Administrator server at a regular interval to check for a work order. When the Avamar Administrator server queues a work order for the client, `avagent` retrieves the work order.

The `avagent` program runs the `avtar` program with the parameters that the work order specifies. The `avtar` program runs the backup based on the set of parameters related to the backup task. The `avtar` program performs a backup by making a connection to the Avamar server over the LAN or a remote connection over the WAN. Avamar uses TCP/IP as the base protocol to make the connection.

Restores are executed in a similar manner to backups. A restore work order is created containing the parameters necessary to complete a restore of all or a subset of the files of a specific backup.

Restrictions and limitations

The following table lists known restrictions and limitations to consider during planning and design. In addition to this table, review the known limitations in the *EMC Avamar Release Notes*.

Table 8 Known restrictions and limitations for planning and designing the Avamar system

Restrictions and limitations	Impact
Recovery time objective (RTO)	RTO involves processes, communication service levels, regular testing, and people. The time to restore data is only one of several critical components needed to achieve a given RTO. Also, the RTO for any individual client is typically limited by the performance capabilities of the client or network, and not the capability of the Avamar server to restore the data.
5 to 10 million files per Avamar client	An Avamar client with several million files can impact backup scheduling. The actual amount of time required to back up the Avamar client depends on the following factors: <ul style="list-style-type: none"> • Total number of files on that client • Hardware performance characteristics of the client The Avamar system can accommodate file system clients with significantly more than 10 million files, but this might require additional configuration or tuning.
500 GB to 2 TB of database data per Avamar client	An Avamar client with large databases to back up can impact backup scheduling. The actual amount of time required to back up the Avamar client depends on the following factors: <ul style="list-style-type: none"> • Total amount of database data on the client

Table 8 Known restrictions and limitations for planning and designing the Avamar system (continued)

Restrictions and limitations	Impact
	<ul style="list-style-type: none"> Hardware performance characteristics of the client <p>The Avamar system can accommodate database clients with more than 2 TB of database data, but this might require additional configuration or tuning.</p> <hr/> <p>Note</p> <p>If the amount of scanned data, for example, new data in the file system plus the size of all the databases combined, is greater than 1 TB, use cache prefixes.</p>
2 to 5 TB of file server data per Avamar client	<p>An Avamar client that is a file server, which protects a large amount of data, can impact backup scheduling. The actual amount of time required to back up the Avamar client depends on the following factors:</p> <ul style="list-style-type: none"> Total number of files on the client Hardware performance characteristics of the client <p>The Avamar system can accommodate clients with significantly more than 5 TB of file system data, but this might require additional configuration or tuning.</p>

Best practices for Avamar clients

- Carefully review the client support matrix with the presales technical engineer:
 - Ensure that Avamar software supports the clients and applications you want to protect.
 - Verify that Avamar software supports all details of the deployment, such as revisions, clusters, third-party plug-ins, and add-ons.
- Consider storing certain data types (very large databases with very high change rates) on Data Domain[®] systems. The following Avamar clients support backup and restore to and from Data Domain systems:
 - Avamar Plug-in for VMware Image Backup
 - Avamar Plug-in for DB2
 - Avamar Plug-in for Lotus Domino
 - Avamar Plug-in for Exchange VSS
 - Avamar Plug-in for Oracle
 - Avamar Plug-in for SAP with Oracle
 - Avamar Plug-in for SQL Server
 - Avamar Plug-in for SharePoint VSS
 - Avamar Plug-in for Sybase ASE
- Consider multi-streaming backups to improve performance for the following Avamar clients:

- Avamar NDMP Accelerator
- Avamar Plug-in for DB2
- Avamar Plug-in for Lotus Domino
- Avamar Plug-in for Exchange VSS
- Avamar Plug-in for Oracle
- Avamar Plug-in for SharePoint VSS

The *EMC Avamar and EMC Data Domain System Integration Guide* provides more information about the use of Data Domain systems as storage for Avamar backups.

Backup window

The backup window is the portion of each day reserved for normal scheduled backups.

- Operational impact—No maintenance activities are performed during the backup window.
- Default settings—The default backup window begins at 8 p.m. local server time and continues uninterrupted for 12 hours until 8 a.m. the following morning.
- Customization—You can customize the backup window start time and duration to meet specific site requirements.

Maintenance window

The maintenance window is the portion of each day reserved for routine server maintenance activities such as checkpoint validation.

- Operational impact—There might be brief periods when backup or administrative activities are not allowed.

Although you can run backups during the maintenance window, doing so impacts both the backup and maintenance activities. For this reason, minimize any backup or administrative activities during the maintenance window. You can perform restores.

Although garbage collection, `hfscheck`, and backups can overlap, doing so might result in I/O resource contention, which can cause both activities to take longer to complete and possibly even to fail.
- Default settings—The default maintenance window begins at 8 a.m. local server time and continues uninterrupted for 12 hours until 8 p.m.
- Customization—Although you cannot directly customize the maintenance window, its start time and duration is derived from backup window settings.

Replication activities

The following activities occur when you replicate data from the local server to a replication target:

- All other maintenance jobs can start.
- All backup work orders are queued immediately.

The following activities occur when the target server is receiving replicated data from the replication source:

- The garbage collection operation cannot start. All other maintenance jobs, such as checkpoint and `hfscheck`, can start.
- All backup work orders are queued immediately.

If WAN throughput causes a bottleneck, overlapping replication with backups is unlikely to affect the amount of time that the replication requires. Additionally, overlapping replication and backups only slightly impacts backup performance.

The following two reasons explain why some clients take a long time to back up:

- WAN bandwidth limits backup throughput for clients. In this case, because the activity level on the Avamar server is relatively low, you can overlap replication with the end of the backup window.
- The clients are large. The time required for a backup is directly proportional to the type and amount of data on the clients being backed up.

Best practices for scheduling client backups

- Minimize the number of groups used to back up clients. Schedule backups during the server's backup window so that backups do not overlap with daily maintenance jobs.
- Use the default maintenance window schedule. Do not deviate from this schedule unless absolutely necessary.
- Set up a separate Avamar server to back up clients when you have a large number of clients that must be backed up outside of the server's backup window.
- Set up multiple Avamar servers to back up clients from around the globe. For example:
 - Set up a server to back up the clients in the Americas.
 - Set up a server to back up the clients in Europe, Middle East, and Africa (EMEA).
 - Set up a server to back up the clients in Asia Pacific and Japan (APJ).
- Limit the amount of time that checkpoint, `hfscheck`, and garbage collection require by carefully managing the capacity on the node:
 - Limit the clients being backed up.
 - Reduce the retention policies.
 - Back up clients with lower daily change rates.
 - Ensure that the garbage collection operation runs every day.
- Limit the amount of time required to perform backups by monitoring the following amounts:
 - Maximum number of files per client.
 - Maximum amount of database data per client.
 - Maximum amount of data per file server.

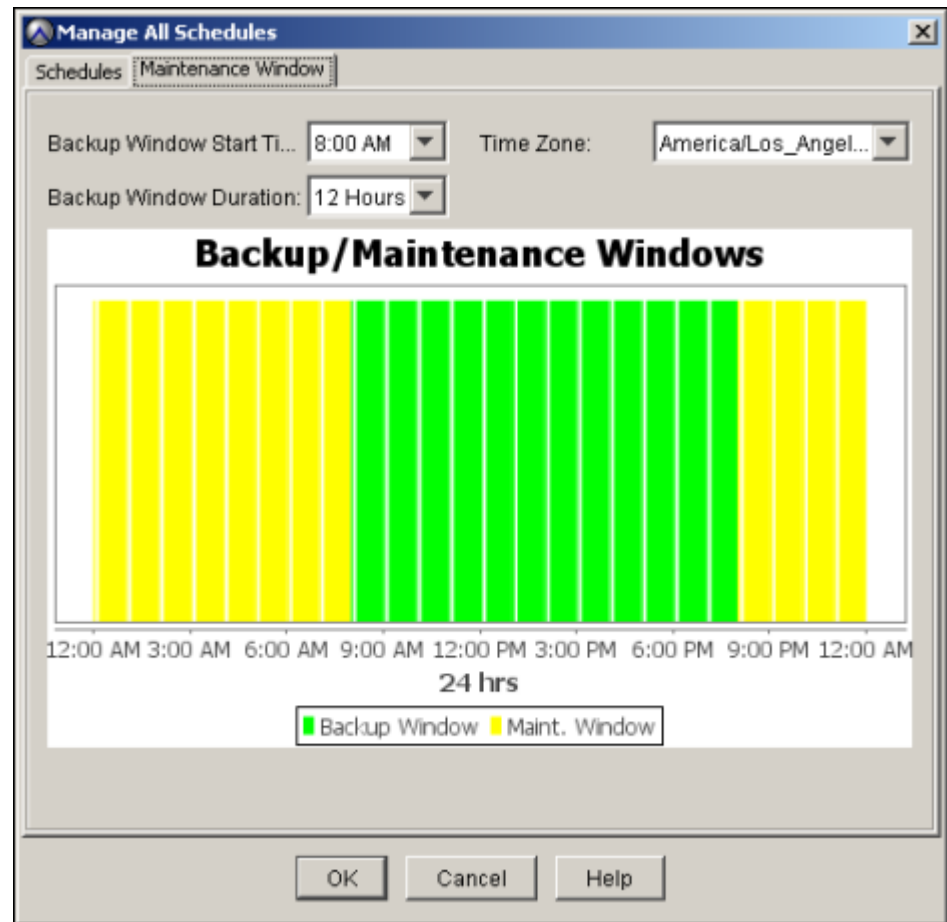
Typically, 80% to 90% of the clients complete daily backups within the first hour or two of the backup window. Consider scheduling replication to start two hours after the start of the backup window. The Avamar server is typically the bottleneck for backup operations only during the first one to two hours of the backup window. The remaining 10% to 20% of the clients might take several hours to complete backups, depending on the number of files or amount of data that needs to be backed up.

Scheduling daily activities

The activities that run the longest during the day typically are `hfscheck`, backups, and replication. During the planning and design stage, ensure to schedule daily activities in a way that best meets system reliability, availability, and supportability.

Each 24-hour day is divided into two operational windows, during which various system activities are performed. The following figure shows the default backup and maintenance windows.

Figure 1 Backup/Maintenance windows



CHAPTER 5

Defining Domains, Groups, and Policies

This chapter includes the following topics:

- [Best practices for making management policy decisions](#)..... 40
- [Defining domains](#)..... 40
- [Defining groups](#)..... 40
- [Defining datasets](#)..... 41
- [Defining schedules and retention policies](#)..... 41

Best practices for making management policy decisions

Make initial backup management policy decisions after you define the overall daily schedule.

- What domains should you set up with designated domain administrators to take advantage of the hierarchical administration capability?
- What groups (which include dataset, backup schedule, and retention policies) should you create to back up clients effectively and manage the backups?
- When should you schedule backups?
- How long should you allow client backups to run?
- How should you set up retention policies to retain the backup data for the required period?

Defining domains

Domains are distinct zones within the Avamar server accounting system that you use to organize and segregate clients. Domains enable hierarchical management. For example, a domain-level administrator can manage all clients and policies within the domain.

Consider segregating clients by domain for billing other internal organizations for backup services. Segregating clients by department or workgroup can be a convenient way to bill them.

If you are not going to use hierarchical management, register all clients in the /clients domain.

Do not add clients to the root (/) domain. Doing so can impact a system migration. To migrate clients from one Avamar server to another, Avamar uses global client IDs (global CIDs). Global CIDs are completely unique Avamar backup client identifiers across all Avamar servers. A system migration does not copy the global CIDs for clients that you locate in the Avamar source system's root domain. If you locate clients in the root domain, you must reregister the clients with the destination server after the system migration completes. The *EMC Avamar Administration Guide* provides more information about system migration.

Best practices for defining domains

- Minimize the number of domains you create. If practical, register all clients in the /clients domain.
- Do not add clients to the root (/) domain.

Defining groups

A group defines the backup policy for the clients that you assign to the group and includes the following three policy elements:

- Dataset policy (including the source data, exclusions, and inclusions)
- Backup schedule (or backup window)
- Retention policy

Best practices for defining groups

- Minimize the number of groups you define. Each dataset policy can include separate dataset definitions for various client plug-ins. For example, a single dataset policy can define independent datasets for Windows, Linux, Solaris, and other clients. You do not need to define separate groups to back up various kinds of operating system clients.
- Leave the default group disabled. By default, all new clients that you activate with the Avamar server are automatically added to the default group. If you enable the default group, any clients that you activate with the Avamar server automatically are backed up according to the default group policy.
- To help manage the capacity on the Avamar server and to avoid being surprised by unexpected clients, leave the default group disabled.
- To back up clients in the default group to the Avamar server, you can add the client to an enabled group and remove the client from the default group.

Defining datasets

In general, do not back up a large client by defining multiple subsets of data that run every night. The use of multiple subsets of data is a good practice in only two instances:

- When you want to define different retention policies for different subsets of data on the client.
- When you break up the dataset so that you back up different subsets of the data on different days of the week.

Best practice for defining datasets

Minimize the number of datasets.

Defining schedules and retention policies

The default schedule runs nightly during the server's backup window. Depending on the amount of data in the largest clients, the default schedule might not be enough time, and you might need to extend the server's backup window. Before you extend the backup window, evaluate the time required for checkpoint, garbage collection, and `hfscheck` to determine that extra time is available after completing these daily activities.

Best practices for defining schedules

- Set appropriate expectations for how long the longest client backups should run every night. Validate that the long-running client backups meet the expectations.
- Minimize the number of clients that you back up outside of the server's backup window.

When you set up backup schedules, remember that mobile laptop clients might require a backup schedule that runs during the day. Laptop clients usually connect to the network during the day. The system can handle a small number of exceptions. In this case, overlap the backups of these laptop clients with the server's maintenance window.

Best practices for setting up retention policies

- Use the advanced retention policy whenever possible to reduce the total amount of back-end storage that the Avamar server consumes:
 - The amount of unique data for weekly backups is equivalent to three daily backups.
 - The amount of unique data for monthly backups is equivalent to six daily backups.

For example, you can configure a retention policy to keep 30 days of daily backups and 3 months of monthly backups. The following figure shows the retention policy settings in the **Edit Advanced Retention Policy** dialog box.

Figure 2 Edit Advanced Retention Policy dialog box



When you use this retention policy, the amount of client data stored on the Avamar is equivalent to the initial unique data plus 42 equivalent days of backups. You use less back-end capacity than the amount you would use if you stored 3 months of daily backups. Three months of daily backups are equivalent to the initial unique data plus 91 equivalent days of backups.

The *EMC Avamar Administration Guide* contains more information about advanced retention policies.

- Set the minimum retention period to at least 14 days.

When you select the maximum retention period, the Avamar server does not retain the last unexpired backup. For a short retention period such as 7 days or less, closely monitor the backup operations to ensure that the last unexpired backup does not expire before the system completes another backup. If all client backups expire before you correct the issue that prevented the client from completing a backup, the next backup is equivalent to an initial backup.

CHAPTER 6

Daily Monitoring of Backup Infrastructure and Operations

This chapter includes the following topics:

- [Monitoring the Avamar system](#)..... 44
- [Monitoring the Avamar system backup operations](#)..... 45

Monitoring the Avamar system

The system reports all Avamar system activity and operational status as events to the administrator server. Examples of Avamar events include offline server nodes, failed or missing server maintenance jobs, and hardware issues.

Monitor the event notification system for warning and error events every day. Monitor the Avamar server daily and understand how to interpret all system warnings and errors.

The following table describes possible ways to monitor Avamar systems.

Table 9 Ways to monitor the Avamar system

Method	Description
syslog or SNMP event notification	If the network management infrastructure supports syslog or SNMP event notification, enable the syslog or SNMP event notification subsystem through Avamar Administrator. The <i>EMC Avamar Administration Guide</i> provides instructions for enabling syslog or SNMP notification.
Email notification system	You can set up email notification to perform the following tasks: <ul style="list-style-type: none"> Batch email notifications that the Avamar system sends twice daily according to the default notification schedule. Send email messages as the selected events occur.
Backup & Recovery Manager user interface	To manually monitor the Avamar system, check the overall health of the Avamar backup infrastructure from the Backup & Recovery Manager user interface.
Unacknowledged events	At least once a day, review and clear any Unacknowledged Events queued: <ol style="list-style-type: none"> From Avamar Administrator, click the Administration launcher button. Select the Event Management tab. Select the Unacknowledged Events tab. <hr/> <p>Note</p> <p>In any Avamar Administrator window, click Have Unacknowledged Events to access the Unacknowledged Events page.</p> <hr/>
Avamar Administrator Event Monitor	At least once a day, review the event monitor: <ol style="list-style-type: none"> From Avamar Administrator, click the Administration launcher button. Select the Event Management tab. Select the Event Monitor tab.

Monitoring the Avamar system backup operations

The system reports all Avamar system activity and operational status as events to the administrator server. You can then use client logs to investigate backup or restore issues. Monitor the event notification system daily for warning and error events related to backup operations. Monitor the **Activity Monitor** daily and understand how to interpret all activity warnings and errors.

Closely monitor daily backup activities

To create consistent backups, you must monitor daily backup activities. The following factors may interfere with backups:

- Network issues
These issues can cause backup failures.
- Client I/O errors
These errors, also known as Completed with Exception status, can prevent the Avamar system from backing up all files.
- High client activity levels
These levels can prevent the Avamar system from backing up all files or can prevent backups from completing within the backup window.
- Operator intervention such as rebooting the client during the backup or canceling the backup
- Incomplete or incorrect dataset definitions
- Inadequate or incorrect retention periods

When you examine the activities, resolve all exceptions and failures.

The most obvious issues are the ones where the clients did not create a restorable backup. The following table describes status messages for these types of failures.

Table 10 Client messages for client backups

Status message	Description
Failed	The client failed to perform the activity. The activity ended due to an error condition. Refer to the associated client log.
Canceled	The activity was canceled, either from the client or from Avamar Administrator. Refer to the associated client log.
Dropped Session	<p>The activity was successfully initiated but, because the Administrator server could not detect any progress, the activity was canceled. The following activities are the most common causes of dropped sessions:</p> <ul style="list-style-type: none"> • Somebody rebooted the client in the middle of the backup. • A network communication outage lasted longer than one hour. <p>The Administrator server automatically queues a rework work order if the client backup fails due to a dropped session.</p>

Table 10 Client messages for client backups (continued)

Status message	Description
Timed Out - Start	The client did not start the activity in the scheduled window. This failure is most likely because the client is not on the network.
Timed Out - End	The client did not complete the activity in the scheduled window. This failure requires special attention because there is a lot of system activity with no restorable backup. In this case, subsequent backups continue to fail with the same status unless some change is made, such as tuning the client caches.

A less obvious failure, but one that still requires attention, is a backup that reports the Completed with Exceptions status. In this case, the backup completed but with errors. The errors are usually due to open files that the Avamar system could not back up. Do not ignore this status. Some missing files, such as .PST files, can be significant.

The primary tool for monitoring daily backups is the **Activity Monitor** in Avamar Administrator. The *EMC Avamar Administration Guide* provides more information about the **Activity Monitor**.

Avamar Administrator can email reports to you that can help you monitor client backups that fail or complete with exceptions. The following table describes these reports.

Table 11 Avamar Administrator reports for client backups

Report	Description
Activities - Exceptions	This report lists all activities for a specific period that completed with exceptions.
Activities - Failed	This report lists all activities for a specific period that failed due to errors.
Clients - No Activities	This report lists all clients with no activities for a specific period.

The *EMC Avamar Reports Guide* provides descriptions for these and other reports that are available.

Best practices for monitoring daily backup activities

- Monitor backups daily and investigate all the failed backups, the missing clients, and the backups that completed with exceptions.
- Enable the advanced statistics report during all backups. This information is useful for addressing performance issues.
- Enable debugging messages when investigating backup or restore failures.
- Enable various activity report email messages, such as:
 - Activities - Exceptions
 - Activities - Failed
 - Clients - No Activities

Closely monitor nightly replication

Ensure that nightly replication successfully completes. The **Activity Monitor** displays a list of all clients that completed replication activities.

CHAPTER 7

Tuning Performance

This chapter includes the following topics:

- [Demand-paging cache](#)50
- [Client caching](#)53
- [Using cacheprefix](#).....54
- [Tuning replicator](#)..... 56
- [Backup and restore performance](#).....56

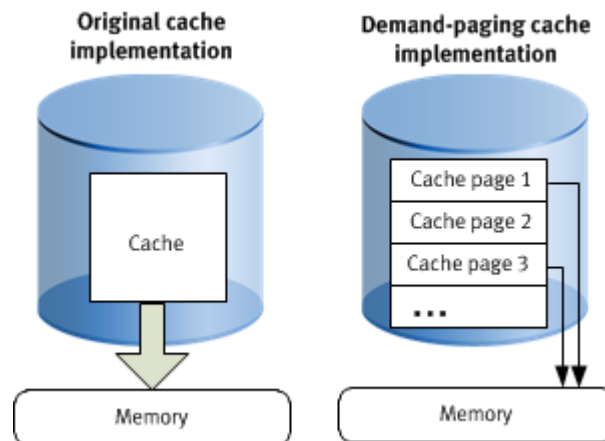
Demand-paging cache

Demand-paging cache is a method for managing Avamar file and hash caches. File and hash caches store information about data that has been backed up to the Avamar server. Starting with Avamar 7.0, demand-paging cache replaces the original Avamar file caching implementation for client backups that you store on a Data Domain system.

Demand-paging cache improves the `avtar` program's capacity to handle a file cache for large file systems by enabling the system to automatically determine the optimal in-RAM cache size for the client environment and to adjust the cache size as needed. Demand-paging cache reduces memory consumption by removing the need to store the entire cache in memory for each instance of `avtar`.

The following figure contrasts the original monolithic cache implementation with the demand-paging cache implementation.

Figure 3 Monolithic cache compared to page cache



The original cache implementation stored file cache information in the `f_cache.dat` file. The demand-paging cache feature uses a new cache file, `f_cache2.dat`, which is stored in the `var` directory. Because the demand-paging cache file uses a unique name, the demand-paging cache feature does not interfere with the original cache file, `f_cache.dat`. The *EMC Avamar Operational Best Practices* for earlier releases of Avamar provides more information about the original monolithic cache implementation.

Demand-paging cache divides the cache into a series of pages. All pages that belong to a backup are kept in a page list. Not all of these pages are memory-resident, but are brought into memory in time order. Demand-paging cache keeps a subsampled list of hashes for each page in memory to handle out-of-order access. This list is called the champion list.

Files that are larger than a threshold size are designated as “must have” entries. These entries almost always are added to the champion list. If a hash is not in the pages in memory, but is in the champion list, the associated page is read into memory. A semi-random selection process designates champions beyond those that are designated as “must have” entries.

File cache

File cache management is implemented by using pages from `f_cache2.dat`. The `avtar` process loads pages, as needed, from `f_cache2.dat` into RAM instead of the entire file.

The original file cache implementation uses approximately 1 GB of disk space and 1 GB of RAM to track approximately 10 million files when `avtar` is running. This data is shared among all of the 16 backups that can be stored in the cache file. This method imposes limits on the maximum size of file systems that Avamar can back up. Demand-paging cache removes this size limitation.

Demand-paging cache requires about 1 GB of disk space to track approximately 10 million files during a backup. The demand-paging file cache can store up to 16 backups worth of nonshared information, which means that the file cache for a normal use case can use approximately 20 GB of disk space. In comparison to the original file cache method, backups that implement the demand-paging file cache require up to 20 times more disk space.

Demand-paging cache, however, enables RAM utilization to remain fairly flat at a much lower size regardless of the number of files in the file system. Automatic tuning within the cache implementation can change the actual RAM utilization amount.

Demand-paging cache example scenario

Each page consists of 100,000 entries and 15,000 champions.

Note

The entries and champion sizes are the maximums per page. The actual utilization is about 60% to 70% of that maximum because of the characteristics of the hash tables, which are used here for their fast access.

For the file cache, each entry is 64 bytes or 6.4 MB per page and each champion is 20 bytes or 300 KB per page.

On disk, the champions are stored with the entries that give a size of about 6.7 MB for each page of the file cache. There is some overhead, but it is a very small percentage of the total size, and usually not relevant.

In RAM, the champions are stored separately. The page sizes are about 6.4 MB for the file cache. The champions are brought into RAM in a single array that contains all of the champions for all of the applicable backups in the cache.

The file cache required for 40 M files is approximately 616 pages: $40 \text{ M} / 65,000$ (estimated) or 4.13 GB on disk ($616 * 6.7 \text{ MB}$).

Auto-tuning alters the number of pages stored in memory based on the size of previous backups and other factors, including the actual memory available to the process. The values may be considerably smaller or larger, based on previous backup requirements.

Cache migration after upgrading the client

After you upgrade a client to Avamar 7.0 or later, the demand-paging cache feature automatically migrates the contents of the `f_cache.dat` file to the new cache file during `avtar` backup operations.

The migration runs when `avtar` determines that demand-paging cache files do not exist or exist but are not valid and when `avtar` finds that `f_cache.dat` files and their relative timestamps fit within the migration period.

The default migration period is 14 days. You can change the default migration period by specifying the `--cache-migration-period` option with the `avtar` command.

Memory improvements from the use of the demand-paging cache feature are not noticeable until after the migration completes. Memory use increases slightly during the migration period since `avtar` uses both forms of cache at the same time.

After the migration period ends, the backup process no longer checks the original cache. After 180 days, the `avagent` process automatically deletes the old cache files. You can change the default time period of 180 days by specifying the `--cachelifetime=days` option with the `avagent` command.

At any time, you can prevent a client migration by removing the `f_cache.dat` cache file before you run an `avtar` backup. You can also prevent a migration by performing one of the following alternatives:

- To prevent a cache migration, run `avtar` with the `--nocache` option. Use of the `--nocache` option stops all caching. Use the `--nocache` option with caution.
- To prevent only a file cache migration, run `avtar` with the `--enable-filecache=false` option.

[Caching options on page 52](#) describes additional options.

Caching options

You can list cache options with the `avtar` command. The following table lists a subset of options that you can specify with the `avtar` command.

Table 12 Options for `avtar` to control demand-paging cache

Option	Description
<code>cache-migration-period=days</code>	Defines the length of time to use the <code>f_cache.dat</code> and <code>p_cache.dat</code> cache files as a source for migration. The default is 14 days.
<code>cacheitmsize=true false</code>	Forces storage of additional access time in the file cache. Use only for the file cache. The default is <code>false</code> .
<code>cachepaxstream=true false</code>	Enables caching for the PAX stream format. The default is <code>false</code> .
<code>cacheprefix=string</code>	Defines a prefix for cache file names. A prefix enables the creation of unique cache files for use with multiple simultaneous instances of <code>avtar</code> .

Table 12 Options for avtar to control demand-paging cache (continued)

Option	Description
	<hr/> Note The <code>--cacheprefix</code> option applies to both sets of cache files: the original cache files and the demand-paging cache files. <hr/>
<code>checkcache</code>	Verifies the integrity of the cache.
<code>clearcache</code>	Deletes existing cache files and starts a new cache file from scratch. This option disables the migration feature from an older-format cache file, if such a file exists.
<code>enable-filecache=true false</code>	Enables or disables the file cache. The default is <code>true</code> .
<code>enable-hashcache=true false</code>	Enables or disables the hash cache. The default is <code>true</code> .
<code>nocache=true false</code>	Enables or disables both file and hash caching. The default is <code>false</code> .
<code>oktoclear</code>	Enables clearing of local cache. The default is <code>true</code> .
<code>paging-cache=true false</code>	Enables or disables the demand-paging cache feature for Avamar backups. The default is setting is <code>false</code> if the target is Avamar Datastore and <code>true</code> if the target is Data Domain. <hr/> Note To revert back to the <code>f_cache.dat</code> cache file, specify <code>false</code> . <hr/>
<code>repaircache</code>	Repairs a corrupted cache file.

Client caching

The `f_cache.dat` and `f_cache2.dat` cache files, which store a 20-byte SHA-1 hash of the file attributes, identify the files previously backed up to the Avamar server. The file cache is one reason why subsequent Avamar backups that occur after the initial backup are fast. When backing up file servers, the file cache screens out approximately 98% of the files. When backing up databases, however, the file cache is not effective because all the files in a database appear to be modified daily.

The `p_cache.dat` file, which is significant when backing up databases, stores the hashes of the chunks and composites that have been sent to the Avamar server. The hash cache identifies the chunks or composites previously backed up to the Avamar server.

The client cache files help reduce the amount of time required to perform a backup and the processing load on the Avamar client and server.

A typical backup should take about one hour for every million files in a file server or about one hour for every 100 GB of data in a database server.

Cache information in the avtar logs

The sizes of the file and hash caches are printed near the beginning of the avtar logs. For example, refer to the following output:

```
avtar Info <5573>: - Loaded cache file C:\Program
Files\Avamar\var\f_cache.dat (5767712 bytes)
avtar Info <5573>: - Loaded cache file C:\Program
Files\Avamar\var\p_cache.dat (25166368 bytes)
```

The file cache is 5.5 MB and the hash cache is 24 MB:

```
1 MB = 1048576 bytes
5767712 bytes/1048576 bytes = 5.5 MB
25166368 bytes/1048576 bytes = 24 MB
```

The end of the avtar log contains the following set of messages:

```
avtar Info <5587>: Updating cache files in C:\Program Files
\Avamar\var
avtar Info <5069>: - Writing cache file C:\Program
Files\Avamar\var\f_cache.dat
avtar Info <5546>: - Cache update complete C:\Program
Files\Avamar\var\f_cache.dat (5.5MB of 63MB max)
avtar Stats <6151>: File cache: 131072 entries, added/updated
140, booted 0
avtar Info <5069>: - Writing cache file C:\Program
Files\Avamar\var\p_cache.dat
avtar Info <5546>: - Cache update complete C:\Program
Files\Avamar\var\p_cache.dat (24.0MB of 31MB max)
avtar Stats <6152>: Hash cache: 1048576 entries, added/updated
1091, booted 0
```

You can see that the file cache has room to increase in size:

```
Files\Avamar\var\f_cache.dat (5.5MB of 63MB max)
```

But the hash cache is at its maximum allowable size:

```
Files\Avamar\var\p_cache.dat (24.0MB of 31MB max)
```

If the file cache is undersized, the “booted” value is nonzero, and the log includes a warning that the cache is undersized. This information is important because the size of the cache influences the overall performance of the system.

Using cacheprefix

When a client does not have enough memory to accommodate the cache files of appropriate sizes, you can back up the client and get the full benefit of appropriately-sized cache files by taking one of the following actions:

- Breaking the client file system into multiple smaller datasets.
- Ensuring that the maximum file and hash caches assign a unique `cacheprefix` attribute for each dataset.

Example for using the cacheprefix attribute

Assume a client has 5.5 million files but only 1.5 GB of RAM.

- One volume has 2.5 million files.
- Three other volumes have 1 million files each.

You can break this client file system into four datasets. A volume with 2.5 million files requires a file cache of at least 110 MB (2.5 x 44 MB). The next increment that accommodates this is 176 MB.

You can define other datasets as shown in the following table.

Table 13 Segregating data into separate datasets

Drive	Attribute settings
C:\ drive (2.5 M files)	<ul style="list-style-type: none"> • <code>filecachemax=220</code> • <code>hashcachemax=3</code> • <code>cacheprefix=driveC</code>
E:\ drive (1.0 M files)	<ul style="list-style-type: none"> • <code>filecachemax=88</code> • <code>hashcachemax=30</code> • <code>cacheprefix=driveE</code>
F:\ drive (1.0 M files)	<ul style="list-style-type: none"> • <code>filecachemax=88</code> • <code>hashcachemax=30</code> • <code>cacheprefix=driveF</code>
G:\ drive (1.0 M files)	<ul style="list-style-type: none"> • <code>filecachemax=88</code> • <code>hashcachemax=30</code> • <code>cacheprefix=driveG</code>

Configure the `cacheprefix` attribute in the dataset by setting the **Enter Attribute** field to `cacheprefix` and the **Enter Attribute Value** field to `driveC`.

The following cache files are located in the Avamar `/var` directory on the client:

```
driveC_f_cache.dat
driveC_p_cache.dat
driveE_f_cache.dat
driveE_p_cache.dat
driveF_f_cache.dat
driveF_p_cache.dat
driveG_f_cache.dat
driveG_p_cache.dat
```

Ensure adequate disk space is available to accommodate the additional file and hash caches.

When specifying various `cacheprefix` values, exclude new cache files from the backups. The cache files are large and have extremely high change rates.

Configuring custom hash settings for Microsoft databases

For a Microsoft Exchange Server or a Microsoft SQL Server database backup, configure the maximum hash cache in the dataset by adding attributes and values.

Procedure

1. From the **Edit Dataset** or **New Dataset** dialog box, select the **Options** tab.
2. Click **More**.
3. In the **Enter Attribute** field, type `[avatar]hashcachemax`.
4. In the **Enter Attribute Value** field, type 200.

Tuning replicator

Work with EMC Customer Service to configure and tune the replicator.

EMC Customer Service performs the following tasks.

Procedure

1. Computes the bandwidth-delay-product (BDP) to determine whether the BDP is high enough to require customized tuning.
2. Verifies that the expected bandwidth is available between the replicator source utility node and the replicator destination storage nodes.
3. Tests the WAN link with the Avamar system components to verify that the Avamar system can utilize about 60% to 80% of the available bandwidth.
4. Sets up the appropriate replication parameters to optimize utilization of the available bandwidth.
5. Tests the replicator to verify its performance.

Backup and restore performance

The backup or restore performance for Avamar 7.2 clients can be impacted when the encryption option is set to high. Unless you require encrypted backups and restores, EMC recommends that you set encryption to none for better performance.

The *EMC Avamar Product Security Guide* provides more information about encryption.

CHAPTER 8

Understanding DPN Summary Reports

This chapter includes the following topics:

• DPN Summary report	58
• Running a DPN Summary report	58
• Example DPN Summary entry	59
• Background on backups	62
• Summary of key DPN summary terms	64

DPN Summary report

Use DPN Summary reports to determine how well an Avamar system performs after it has achieved steady state.

The DPN Summary report helps you to determine the following details:

- Daily change rate for each individual client
- Daily change rate across the overall system
- High change rate clients that contribute the most to overall system change rate
- Amount of data that the Avamar system protects per client and across the system
- Number of clients that the Avamar system protects
- Abnormal client behavior such as:
 - Days with unusually high change rates
 - Unusually long backups
 - Frequent backup failures
- Amount of data that moves across the network with Avamar instead of incremental tape backups
- Benefits from the combined effect of commonality factoring and compression, when compared with commonality factoring or just compression

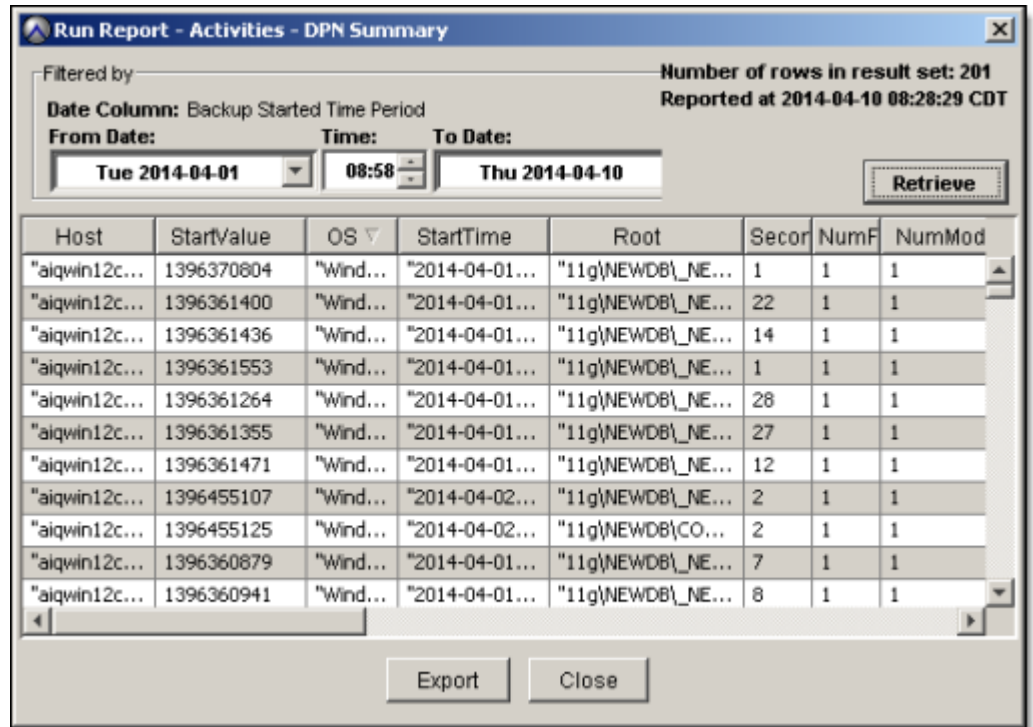
Running a DPN Summary report

You can access the DPN Summary report from Avamar Administrator.

Procedure

1. In Avamar Administrator, select **Tools > Manage Reports**.
The **Manage All Reports** window appears.
2. Select **Activities - DPN Summary** from the navigation tree and click **Run**.
3. Select a date range and click **Retrieve**.

The **Run Report - Activities - DPN Summary** window appears.



4. Click **Export** to save the report to a comma delimited (.csv) file.

Example DPN Summary entry

The DPN Summary report provides statistics for server activities.

The information in the following table uses values from a sample DPN Summary report to provide a scenario for interpreting server activity for a scheduled backup. [Background on backups on page 62](#) provides more information.

Table 14 DPN Summary column descriptions

Column	Value	Description
Host	avamar1.example.com	<p>The client hostname as defined in DNS.</p> <ul style="list-style-type: none"> During backups, the hostname is the client that backs up data to the Avamar server. During restores, the hostname is the client that receives the restored data. <p>Note</p> <p>This client is not the one that sourced the data.</p>
StartValue	1169366400	The UNIX start time of the activity. The UNIX start time is in the local time of the Avamar server.
OS	Windows Server 2008 RS	The client operating system.
StartTime	2013-09-18 09:40:39.36	The date and time the activity starts. The StartTime is in Coordinated Universal Time (UTC)/Greenwich Mean Time (GMT).

Table 14 DPN Summary column descriptions (continued)

Column	Value	Description
Root	/EMC IT Windows Dataset	The name of the dataset that the activity uses, if applicable.
Seconds	2,777	The duration, in seconds, of the activity.
NumFiles	517,023	The total number of files scanned during the activity less those files that were excluded through exclusion rules.
NumModFiles	1,908	The total number of modified files associated with the activity.
ModReduced	55,023,086,382	The amount of modified data that is reduced due to compression during commonality processing.
ModNotSent	4,833,745,400	The amount of bytes in modified files that do not have to be sent to the Avamar server because of subfile-level commonality factoring.
ModSent	4,833,745,410	The amount of new bytes sent to the Avamar server.
TotalBytes	451,940,965,688	Summary of key DPN summary terms on page 64 provides a description for TotalBytes
PcntCommon	99	Commonality percentage during the activity.
Overhead	60,055,981	<p>The number of bytes for COMPOSITES and DIRELEMS used to store data. Overhead is the amount of nonfile data that the client sends to the server for the following items:</p> <ul style="list-style-type: none"> • Indexing information • Requests from the client to the server for the presence of specific data chunks • ACLs • Directory information • Message headers <p>On any active file system, overhead is usually a small percentage of the file data that is sent to the Avamar server.</p>
WorkOrderID	EMC IT Windows Schedule— 1169348400105	<p>The unique identifier for the following activities:</p> <ul style="list-style-type: none"> • For scheduled backups, the format of a work order ID is <i>schedule_name-group_name-unix_time</i> in milliseconds where <i>schedule_name</i> is the name of the Avamar schedule and <i>group_name</i> is the name of the Avamar group. • For on-demand backups you start by selecting the Back Up Group Now option in the Policy window, the format of the work order ID is <i>group_name-unix_time</i> in milliseconds. • For on-demand backups or restores you start by selecting the Backup and Restore option from

Table 14 DPN Summary column descriptions (continued)

Column	Value	Description
		<p>the Policy window, the format of the work order ID is <i>mod-unix_time</i> in milliseconds.</p> <ul style="list-style-type: none"> For on-demand backups you start from the systray icon on a Windows Avamar client, the format of the work order ID is <i>cod-unix_time</i> in milliseconds. For command-line backups or restores, the format of the work order ID is <i>nah-unix_time</i> in milliseconds. For replication activities, the format of the work order ID <i>cod-nah-unix_time</i> in milliseconds.
ClientVer	7.1.100-135	The Avamar client software version.
Operation	Scheduled Backup	<p>Operation is one of the following types of activities:</p> <ul style="list-style-type: none"> On-demand backup Scheduled backup Restore Validate Replication source Replication destination
Status	Activity completed successfully.	<p>The FINAL status of the client activity is one of the following types:</p> <ul style="list-style-type: none"> Activity completed successfully Activity completed with exceptions Activity cancelled Activity failed - timed out before starting Activity failed - timed out before completion Activity failed - client was given a workorder, but did not acknowledge its receipt Activity failed - client error(s) Activity failed - timed out before completion Activity failed - client has no data specified by dataset Dropped Session - No progress reported
SessionID	9116934840011000	The SessionID is a unique identifier for the client activity.

Background on backups

This topic provides background information about how the Avamar client performs backups, including key statistics.

Dataset size

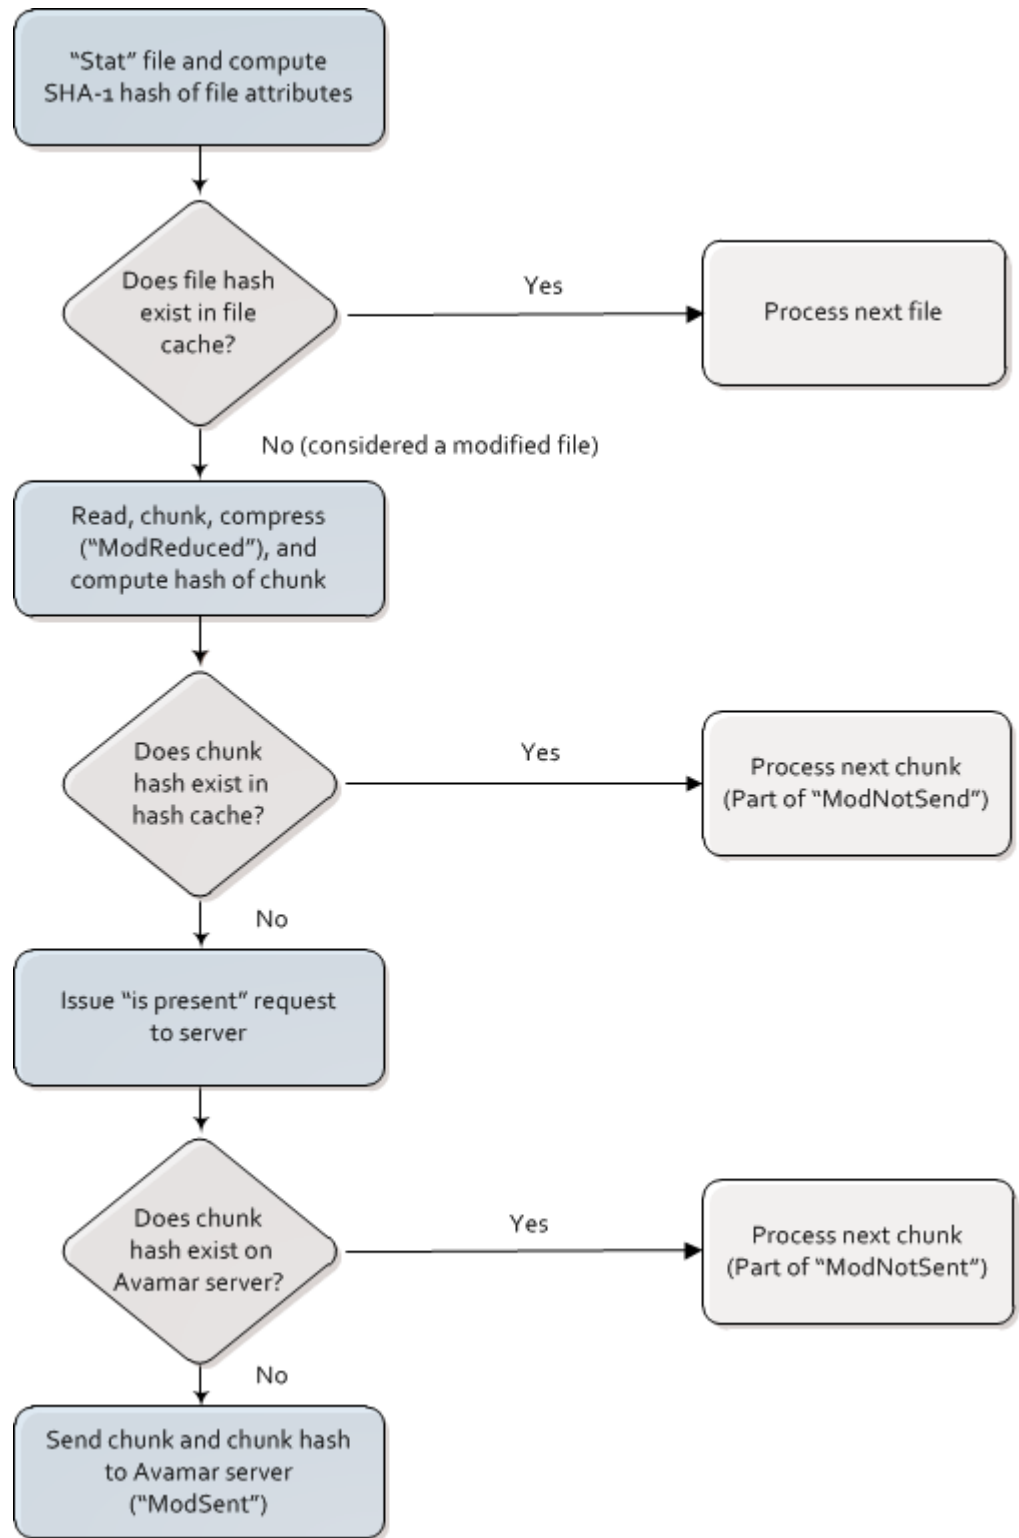
Begin with the value in the TotalBytes column as shown in [Table 14 on page 59](#). This value, 451,940,965,688 bytes (or 421 GB), represents the dataset size. This total does include files that you exclude with exclusion rules or open files that you could back up, perhaps because the file system was not frozen.

Modified files

When scanning through a file system, obtain file metadata and compute the SHA-1 hash of the metadata. Then look up SHA-1 hash in the file cache on the client. If the hash is present, the opening and reading of the files is not necessary. Therefore, a high percentage of hits in the file cache makes the overall backup proceed quickly.

Any file whose metadata hash gets a miss in the file cache is a modified file. A file that was modified since the last backup. Therefore, the Mod bytes in NumModFiles, ModReduced, ModNotSent, and ModSent columns are really shorthand for bytes associated with modified files. For example, files you must open and read so that all the data in the file can be chunked, compressed, and hashed.

The following figure shows the roles of the file cache, the hash cache, and the server is_present requests in determining which data to send to the server.

Figure 4 Data flow to server

The cache and hash flowchart references the following terms:

- ModReduced

When Avamar backs up modified files, the data is chunked, compressed, and then hashed. Because the compression takes place on the client, the amount of

compressed data is reported as ModReduced. In [Table 14 on page 59](#), the ModReduced = 55,023,086,382 (51 GB).

- ModNotSent

When subfile level commonality exists, the data is not sent. ModNotSent is shorthand for bytes in modified files that do not have to be sent to the Avamar server because of subfile-level commonality factoring. ModNotSent = 4,115,205,370 (3.8 GB) in [Table 14 on page 59](#) means 3.8 GB of compressed chunks were already on the Avamar server.

- ModSent

When new bytes are sent to the server, they are reported as ModSent. In this case, ModSent = 393,498,485 (0.37 GB).

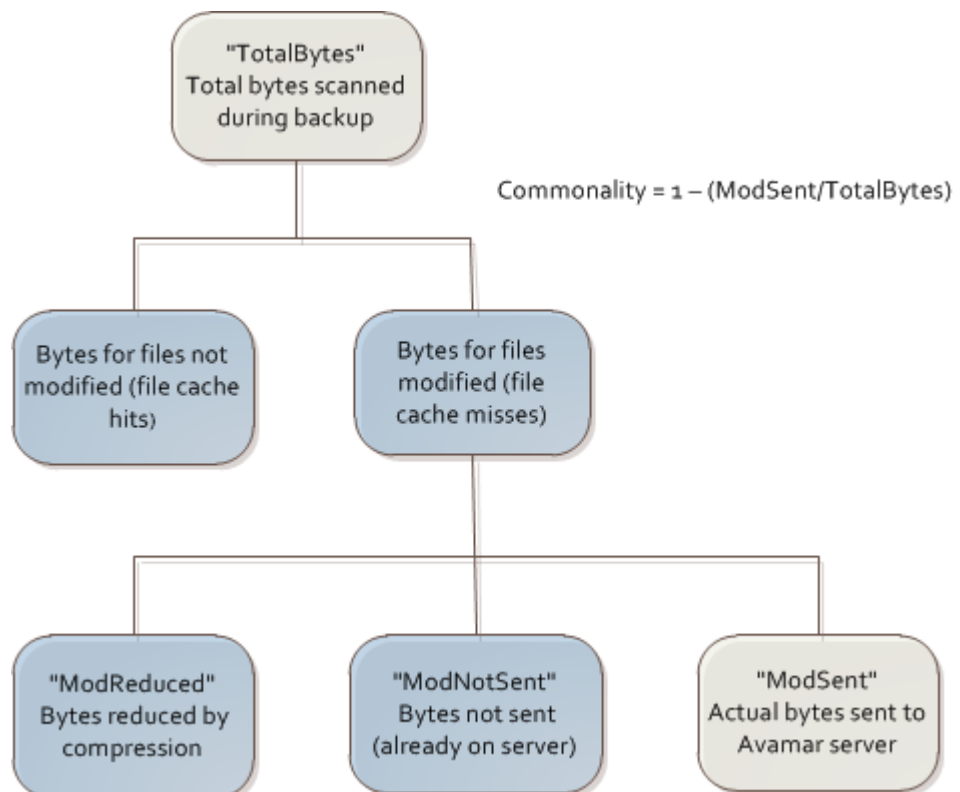
Summary of key DPN summary terms

The following describes the relationships between DPN summary terms:

- TotalBytes = (Total bytes in the dataset, including open files that were not backed up) - (Subtotal bytes excluded by exclusion rules)
- TotalBytes = (Subtotal bytes for files not modified) + (Subtotal bytes for files modified since previous backup)
- Subtotal bytes for files modified since previous backup = ModReduced + ModNotSent + ModSent

The following figure shows the relationship between these values.

Figure 5 Avamar commonality diagram



Definition of commonality

Avamar change rate is equivalent to the ModSent divided by TotalBytes: $\text{ModSent} / \text{TotalBytes}$. The Avamar commonality equals 1 minus the change rate.

Avamar backup compared to incremental tape backup

During an incremental tape backup, the amount of data that the Avamar system sends across the network (if the backup data is not compressed on the client) is equal to the Subtotal bytes for files modified since the previous backup.

The following formula calculates the efficiency of the Avamar commonality factoring when compared to incremental tape backups:

$\text{ModSent} / (\text{ModReduced} + \text{ModNotSent} + \text{ModSent})$

On the particular date shown in [Table 14 on page 59](#), the Subtotal bytes for files modified since the previous backup = 51 GB + 3.8 GB + 0.37 GB = 55 GB. If you divide this amount by the TotalBytes, the result is 55/421 or 13%.

Usually in day-to-day backups of file servers, you expect this value to be in the 2% range.

When backing up databases, expect this value to be 100% because every file is touched every day. The total bytes for modified files, therefore, is equal to the total bytes.

Definition of terms during restore activities

During restore and validate activities, ModReduced is the amount that the data expands during the restore or validate operation. ModSent is the amount of data the Avamar server sends to the Avamar client during the restore or validate operation. During restore or validate, $\text{TotalBytes} = \text{ModSent} + \text{ModReduced}$.

CHAPTER 9

Protecting Avamar Desktop/Laptop Clients

This chapter includes the following topics:

• About Avamar Desktop/Laptop	68
• Deploy additional Avamar servers for Desktop/Laptop clients	68
• Create a dataset to back up only user data	68
• Keep the initial client backups to a manageable number	71
• Consider node size when configuring Avamar servers	73
• Determine the backup window	73
• Schedule the backup window	74
• Adjust runtime for daily maintenance tasks	74
• Do not run client utilities during the backup window	75
• Run backups more frequently than the retention policy	75
• Prevent backups on a wireless connection	75
• Manage storage capacity for Desktop/Laptop clients	76
• Ensure adequate initialization time for Wake-on-Lan backups	76

About Avamar Desktop/Laptop

Avamar Desktop/Laptop is client/server software that extends data backup and recovery to users who are on the LAN, in remote offices, or connected to the corporate network through a VPN.

When users log in during normal backup windows, Avamar Desktop/Laptop backs up data from the desktop and laptop computers to the Avamar server by using existing network links. Users can also start backups from the desktop user interface.

You install the Avamar Desktop/Laptop client as part of the Avamar Client for Windows, the Avamar Client for Mac OS X, or the Avamar Client for Linux installation. You install the Avamar Desktop/Laptop server as part of every Avamar server installation.

The *EMC Avamar Administration Guide* provides more information about Avamar Desktop/Laptop.

Deploy additional Avamar servers for Desktop/Laptop clients

When deploying Avamar Desktop/Laptop to a location with existing Avamar servers, use an additional Avamar server to support the desktop and laptop clients. You must run backups of Desktop/Laptop clients when users are online (usually during the day). [Adjust runtime for daily maintenance tasks on page 74](#) provides more information. Scheduled backups for file servers and database clients normally run during the night. [Scheduling daily activities on page 38](#) contains more information about backup and maintenance windows.

Best practice for deploying Desktop/Laptop clients

When deploying Avamar Desktop/Laptop to a location with existing Avamar servers, use an additional Avamar server to support the desktop and laptop clients.

Create a dataset to back up only user data

The use of Avamar Desktop/Laptop to back up users' desktop and laptop computers can impact Avamar storage capacity depending on the following factors:

- Number of desktop and laptop computers to back up
- Amount of data on each computer

To best manage storage capacity, back up only the user files and folders, and exclude the common data such as application and operating system files.

Best practices for creating a dataset

- Create a backup dataset that specifies the files and folders for the backup.
- Exclude certain file types from desktop and laptop backups.
- If practicable, minimize the number of entries you define in exclude and include lists.
- In an environment that contains both Windows XP and Windows Vista or Windows 7 clients, add the `--x18=256` option to the dataset to prevent the `Path not found` error.

The following table lists folders to include in a Desktop/Laptop dataset.

Table 15 Desktop/Laptop file types to include in a dataset

OS	Tab	Files and folders
Windows	Source Data	<ul style="list-style-type: none"> • <code>#USERDOCS#*\Desktop</code> • <code>#USERDOCS#*\Documents</code> • <code>#USERDOCS#*\My Documents</code> • <code>#USERDOCS#*\Favorites</code> <hr/> <p>Note</p> <p>A change in the default location of the user directories was made between the Windows XP release, and the Windows Vista and Windows 7 releases. To handle this change, the Windows Desktop/Laptop plug-in uses <code>#USERDOCS#</code> as a variable that translates to the default location based on the specific Windows operating system.</p>
Mac	Source Data	<ul style="list-style-type: none"> • <code>/Users</code> • <code>/Users/*/Desktop</code> • <code>/Users/*/Documents</code> • <code>/Users/*/Library/Safari</code>

Exclusions for Windows computers

The following table lists folders to exclude from the Desktop/Laptop dataset for a Windows system.

Table 16 Desktop/Laptop file types to exclude from a dataset

File type	Files and folders
Link files in each user's Recent folder	<code>*\Recent*.lnk</code>
Google Desktop Search folder	<code>*\Local Settings\Application Data\Google\Google Desktop Search</code>
Windows Indexing and Search services	<ul style="list-style-type: none"> • <code>*\catalog.wci</code> • <code>*\windows.edb</code>
Other nonbusiness files such as personal music, pictures, video, and so forth	<ul style="list-style-type: none"> • <code>*.avi</code> • <code>*.cdr</code> • <code>*.dmg</code> • <code>*.iso</code> • <code>*.m4v</code> • <code>*.mov</code> • <code>*.mp3</code> • <code>*.mp4</code> • <code>*.mpeg</code>

Table 16 Desktop/Laptop file types to exclude from a dataset (continued)

File type	Files and folders
	<ul style="list-style-type: none"> • *.jpeg • *.rar • *.r[0-9][0-9] • *.tgz • */iTunes/ • *.wma • *.wmv
Anti-virus software quarantine files	Check the antivirus vendor documentation to determine the folder used to store quarantine files.
Recycle bin files	<p><SYSTEM DRIVE>\%Recycle.bin (Windows Vista and Windows 7)</p> <p><SYSTEM DRIVE>:\RECYCLER (Windows XP)</p>

Exclusions for Mac computers

Exclude the following files and folders in the Desktop/Laptop dataset for a Mac system:

- */.Trash/
- */Library/Caches/
- */Library/Cookies/
- */Library/Logs/
- */Library/PubSub/Feeds/
- */Library/Application Support/SyncServices/Local/

Minimize the number of exclude and include lists

Avamar must compare every file you select for a backup with each entry in both lists to determine whether to back up the file. This comparison process adds overhead and potentially increases the duration of each backup.

Dataset caveat

In an environment that includes desktop or laptop computers on Windows XP and on Windows Vista or Windows 7, backups can appear to fail if you use a single dataset that specifies the `My Documents` folder and the `Documents` folder.

A backup in such an environment displays a backup failure on the status bar and writes an error similar to the following to the log file:

```
Path not found
```

In an environment that contains both Windows XP and Windows Vista or Windows 7 clients, add the `--x18=256` option to the dataset to prevent the `Path not found` error.

Procedure

1. In Avamar Administrator, select **Tools > Manage Datasets**.

The **Manage All Datasets** dialog box appears.

2. Select the dataset from the list and click **Edit**.

The **Edit Dataset** dialog box appears.

3. Click **Options** and select the plug-in from the **Select Plug-In Type** list.

4. Click **More**.

5. Type **x18** in the **Enter Attribute** text box.

6. Type **256** in the **Enter Attribute Value** text box.

7. Click **+**.

The attribute/value pair (**--x18=256**) appears in the large text box below the **+** and **-** buttons.

8. Click **OK** twice.

Keep the initial client backups to a manageable number

Avamar Desktop/Laptop environments can support up to 5000 clients for each Avamar server. However, simultaneously running first-time backups for hundreds or thousands of clients can create network throughput issues. These issues can prevent the completion of the backups within the backup window. Throughput issues caused by large amounts of data transfer are normally only an issue when running first-time backups. The savings you realize through data deduplication are at their lowest when you first back up a system. At this time the amount of data that must be transferred is at its greatest.

Best practices for first-time backups

- Keep the initial client backups to a manageable number.
- To minimize the impact of first-time backups, bring clients online in smaller groups.

Strategy for performing first-time backups

All new Avamar client computers require a first-time backup of all the data that the dataset specifies. These first-time backups require more time and storage space than subsequent backups that only back up changed data.

Use the first few groups to discover information about the capabilities of the network. Start with smaller groups of clients and, after you successfully add each group, increase the size of the next group.

On the first day, start with activation and a first-time backup of clients equal to no more than 10 times the number of storage nodes deployed. For example, if you have 5 storage nodes, back up 50 clients:

$$10 \times 5 \text{ storage nodes} = 50 \text{ backup clients}$$

If all backups for the first day complete within the scheduled backup window, double the amount of clients in the group on day two. Continue adding more clients on subsequent days until all initial backups are complete. Reduce the number of clients if any backups fail to complete within the backup window.

Strategy for setting up backup groups

As the amount of data archived in an Avamar system increases, the benefit of global deduplication increases. This means that throughput problems decrease exponentially as you add more clients and the number of common data objects in the system increases.

Procedure

1. Place clients with the smallest burden on the network infrastructure in the first backup groups that you bring online.
2. Place clients with the greatest burden on the network in the last groups brought online.
3. Set up backup groups following these guidelines:
 - a. Ensure that the first backup groups consist of computers that traverse the shortest logical distance to the Avamar server.

The following factors increase the logical distance:

- Routers
- Firewalls
- VPN tunnels
- Physical distance

- b. Ensure that first backup groups consist of computers that use the fastest network transmission rates.

An illustrative, nonexhaustive list, from fastest to slowest:

- Gigabit Ethernet
- 802.11n
- Fast Ethernet
- 802.11g
- Ethernet
- V.92
- V.34

Activating clients by using the Avamar Client Manager

You can bring clients online in sequentially targeted groups by using the directory information tree or the search capability of Avamar Client Manager.

Procedure

1. You can select appropriately sized and situated organizational units by using the tree structure. Or, you can use search terms to define a group with the target number and type of clients.
2. Then, use the drag and drop capability to activate and initiate first-time backups of correctly sized and connected groups of clients.

Results

This way, you avoid the problems associated with overtaxing the throughput capabilities of the network.

Consider node size when configuring Avamar servers

In certain situations, the use of lower capacity nodes can be advisable. The use of lower capacity nodes increases the connection count, which in turn, increases the potential number of concurrent backups.

Consider node size when configuring Avamar servers.

Determine the backup window

Avamar Desktop/Laptop environments include more clients than traditional Avamar systems. The Avamar Administrator server allows a maximum of 72 concurrent backup connections for each active storage node. The Avamar server reserves one connection for restores.

To determine how many backups can complete within an hour, consider the following two examples.

Example 1 10 minutes per backup

Backup criteria:

4 storage nodes

Average backup time = 10 minutes

6 backups per storage node connection per hour (60 min./10 min. = 6)

Formula:

4 nodes x 6 backups per connection x 72 concurrent connections = 1,728 backups per hour

Example 2 40 minutes per backup

Backup criteria:

5 storage nodes

Average backup time = 30 minutes

2 backups per storage node connection per hour (60 min./30 min. = 2)

Formula:

5 nodes x 2 backups per connection x 72 concurrent connections = 720 backups per hour

In both examples, the following variables can affect the total backup time:

- Total size of the backup dataset
- Amount of daily changes to data
- Network speed and amount of network traffic

Example 2 40 minutes per backup (continued)

- Concurrent processes that run on the Avamar server

For instance, backing up data from a LAN-connected laptop usually takes less time to complete than backing up the same computer when it is connected to the corporate network by an 802.11G wireless connection.

Best practice for determining the backup window

Use the number of backups per hour to help you determine a backup window, which allows enough time to back up all desktop and laptop computers.

Schedule the backup window

The backup window for desktop and laptop clients is often the opposite of traditional server clients. Avamar Desktop/Laptop backups must run while the desktop and laptop computers are online. The backup window typically is during the work day.

When determining the backup window, ensure that it is flexible enough for users who are offline due to traveling, meetings, and so forth.

Start with a backup window of 12 hours and increase or decrease it as necessary. Depending on the location of remote clients, backing up all clients can require multiple Avamar servers.

Best practice for scheduling backups

Schedule the backup window to back up desktop and laptop computers when they are most likely to be online.

Adjust runtime for daily maintenance tasks

It is important that Avamar daily maintenance tasks complete successfully every day. Failures of these tasks quickly result in capacity problems for the Avamar server.

The timing for daily maintenance tasks for Avamar Desktop/Laptop is different from the timing of a standard servers-as-clients deployment of Avamar.

In a standard deployment, the backups of servers occur at night when they are least active. To accommodate this, the Avamar daily maintenance tasks run during the day.

For Avamar Desktop/Laptop, the client backups usually run during the day, when the clients are most likely to be powered-up and connected. You often must change the Avamar daily maintenance tasks to run during the night to avoid conflicts with client backups and restores. [Scheduling daily activities on page 38](#) provides more information.

The daily maintenance task of garbage collection requires a quiet system. Garbage collection cannot start when backups are running. Because garbage collection reclaims space on the Avamar server, the failure to successfully complete this task can quickly create capacity problems.

Best practice for scheduling daily maintenance tasks

Adjust the maintenance window so that it allows daily maintenance tasks to complete without overlapping with the backup window.

Do not run client utilities during the backup window

Avoid running multiple system utilities on the user's PC or Mac simultaneously with backups. For instance, do not schedule antivirus scans or disk defragmenter jobs during the backup window.

Run backups more frequently than the retention policy

Backup retention policies specify how long to keep a particular backup in the system. The Avamar system automatically marks backups for deletion after they expire.

Run backups more frequently than the amount of time you specify by the retention policy. Frequent backups ensure that data is always available from the backup. For example, if the retention policy is 30 days, ensure that you run backups before the 30-day retention period expires. If you fail to back up data within the retention period, the data is no longer part of the backup. The data is still available on the hard drive unless it has been deleted.

[Defining Domains, Groups, and Policies on page 39](#) and *EMC Avamar Administration Guide* provide more information about retention policies.

As an alternative to running backups more frequently than what the retention policy specifies, you can configure the Avamar server to never delete the client's last backup. To enable this feature you set the `keep_last_backup` key to `true` in the `mcserver.xml` file on the Avamar server. This setting prevents Avamar from deleting a client's last backup, which can be beneficial for backup clients with short retention policies (less than 2 weeks), or backup clients that are offline for extended periods of time when users are out of the workplace for vacations or for other reasons. Configuring the Avamar server to never delete the client's last backup can be preferable to running backups more frequently than what the retention policy specifies. The tradeoff in configuring the `keep_last_backup` key in the `mcserver.xml` file on the Avamar server is that more space is consumed on the Avamar server because the last backup for all clients is never deleted.

Note

When you enable the `keep_last_backup` key, this setting affects all clients, including all non-Desktop/Laptop clients. For example, if a client runs multiple backups with different plug-ins, only one of these backups is retained.

The *EMC Avamar Administration Guide* provides more information about setting the `keep_last_backup` key.

Prevent backups on a wireless connection

In some locations, users can pay exorbitant data transmission fees when their backups run over a wireless connection. To avoid these exorbitant data transmission fees, you can disable backups from running over a wireless connection by clearing the **Back Up On Wireless** option.

Clearing the Back Up On Wireless option on Windows

Procedure

1. Right-click the Avamar icon in the system tray.
The context menu appears.
2. Select the **Settings** menu.
3. Clear the **Back Up On Wireless** option.

Clearing the Back Up On Wireless option on Mac

Procedure

1. Click the Avamar icon on the menu bar.
The context menu appears.
2. Select the **Settings** menu.
3. Clear the **Back Up On Wireless** option.

Manage storage capacity for Desktop/Laptop clients

The most important consideration in successfully maintaining Avamar Desktop/ Laptop is capacity management. An appropriately managed Avamar server achieves steady state when storage capacity is well below the capacity warning threshold, which is 80% of storage capacity.

You achieve steady state operation when the average data sent to the Avamar server is less than or equal to the average data removed from the multi-node server. [Steady state system on page 31](#) provides more information.

As a multi-node server enters the range of 85% to 100% of storage capacity, performance degradation occurs. If storage capacity reaches 100%, the Avamar server transitions to a read-only state. This transition protects the integrity of the data already stored on the server. [Avamar capacity thresholds on page 28](#) provides more information.

A server cannot achieve steady state and will exceed storage capacity for the following circumstances:

- Clients back up more than the initial backup size limit
- Clients exceed the daily data change rate limit
- Garbage collection fails

Manage storage capacity for Avamar Desktop/ Laptop clients.

The *EMC Avamar Administration Guide* provides more information about capacity management. This guide provides detailed descriptions of the features, functions, and tools available to assist you with correctly monitoring and managing server storage capacity.

Ensure adequate initialization time for Wake-on-Lan backups

Both Windows and Mac OS X offer power management features to reduce power consumption and save energy. If you use Wake-on-Lan (WoL) network technology to

remotely power on or wake up a computer before a scheduled backup starts, ensure client systems have adequate initialization time.

Best practices for Wake-on-Lan backups

- Ensure that power management settings for client computers do not return the client to a powered-down or sleep state before the Avamar system receives the backup request.
- Depending on the number of clients you back up, the Avamar system may queue clients while waiting for processing.
- Schedule WoL backups so that clients are powered on or awake before a connection is available.

CHAPTER 10

Other Avamar Administration Best Practices

This chapter includes the following topics:

- [Protecting the Avamar server](#).....80
- [Changing passwords](#).....80
- [Using Avamar Client Manager](#).....81
- [Enabling the Email Home feature](#).....82
- [Using EMC Secure Remote Support solution](#).....82
- [Assigning users](#).....82

Protecting the Avamar server

Deploy an Avamar server in a protected network and not one exposed to the Internet. Even when you deploy an Avamar server in an internal network, protect the server by a firewall to prevent unauthorized access to the nodes that comprise the server. Protect the Avamar server from the Internet by providing full firewall protection.

Use of an uninterruptible power supply with Avamar

Under some circumstances, an unclean shutdown of an Avamar server can result in data inconsistencies, which a rollback can recover from, but results in the loss of backups that complete after the last successful checkpoint. There are various ways Avamar servers can experience unclean shutdowns, many of which are preventable. They include the following cases:

- Unexpected site-wide power outages.
- Not performing a clean shutdown before planned power outages.
- Not connecting the Avamar Data Store redundant power system to independent power sources, or using incorrect shutdown procedures.

Avamar systems are not synchronously protected against unexpected power loss. Avamar nodes have dual power supplies, so use an uninterruptible power supply (UPS) to achieve full protection against power loss. EMC recommends the use of redundant power supplies to the Avamar systems that are connected to separate sources, with one of them being protected with a UPS.

After a power outage has occurred and while the system is being protected by the UPS, run a checkpoint in preparation for shutting down the system. Use the proper shutdown procedure from the Knowledgebase article esg112243, “Avamar Data Store Single- and Multi-node Shutdown/Startup Procedures.” This article is available from [EMC Online Support](#).

Best practices for using a UPS with Avamar

- Connect redundant power to separate sources with one of the sources being backed by UPS.
- Use proper shutdown procedures.
- Ensure daily integrity checks are successful, should a rollback be required.
- Deploy a UPS for the Avamar server hardware to protect against data loss caused by unplanned power outages.

Changing passwords

If you have not changed Avamar passwords from the factory default values, use the `change-passwords` utility to change them.

The following table lists Avamar user accounts and SSH keys that require password changes.

Table 17 Avamar user accounts and SSH keys

Type	Username
Operating system user accounts	<ul style="list-style-type: none"> • root • admin • dpn
SSH keys	<ul style="list-style-type: none"> • admin_key • dpnid
Root-level software application user accounts	<ul style="list-style-type: none"> • root • MCUser • replonly

Change the passwords for all these user accounts. Changing only the passwords for the operating system users does not sufficiently prevent someone from logging in to Avamar server nodes. If you have not changed the two SSH keys, someone could use the factory-default SSH keys to log in to the Avamar server. The *EMC Avamar Administration Guide* provides more information about changing Avamar passwords.

Best practice for default passwords

Change all factory default passwords except the passwords for the backuponly, restoreonly, and backuprestore software application users.

Using Avamar Client Manager

Avamar Client Manager is a web-based management application that provides centralized Avamar client administration capabilities for larger businesses and enterprises. You start Avamar Client Manager from the Backup & Recovery Manager user interface by using the launch feature.

For environments that include many clients, use Avamar Client Manager to simplify the following tasks:

- Activating clients
- Moving clients to new domain on same Avamar server
- Removing a client from a group on same Avamar server
- Moving a client to a new group on same Avamar server
- Moving client to new Avamar server
- Reporting with backup and restore summary
- Retiring clients
- Deleting clients
- Upgrading client software (requires the 7.2 plug-in)

Best practice for using Avamar Client Manager

Use Avamar Client Manager to help manage large numbers of Avamar clients. The *EMC Avamar Administration Guide* provides more information about Avamar Client Manager.

Enabling the Email Home feature

When configured and enabled, the Email Home feature, including ConnectEMC, automatically emails configuration, capacity, and general system information to EMC Customer Service once daily, and critical alerts in near-real time on an as-needed basis.

Enable this feature on all Avamar servers. The *EMC Avamar Administration Guide* provides more information. The Email Home feature is offered as part of the server maintenance plan. However, it is offered with the following understanding:

- There is no guaranteed service level agreement for monitoring Email Home messages. You must assume primary responsibility for monitoring the Avamar systems.
- Support cases are automatically opened for issues that affect the backup infrastructure (Avamar server) such as a failed `hfscheck`. Support cases are not opened for issues with backup operations such as failed backups.
- EMC Customer Service will not proactively alert you of problems, such as a server down or a disabled schedule, that prevent Email Home messages from being sent.

Using EMC Secure Remote Support solution

EMC Secure Remote Support Gateway (ESRS) is an IP-based automated connect home and remote support solution. The ESRS IP solution creates both a unified architecture and a common point of access for remote support activities performed on EMC products.

To simplify remote support of Avamar servers, install the EMC Secure Remote Gateway. It is integrated with ConnectEMC.

The use of the ESRS IP solution enables EMC Customer Service to perform the following functions:

- Log in through the EMC Secure Remote Gateway to troubleshoot problems, which eliminates the need for WebEX sessions.
- Begin work on critical issues soon after a ConnectEMC notification of a problem is received.

Best practice for EMC Secure Remote Support

Enable ESRS.

The *EMC Secure Remote Support IP Solution Security Management and Certificate Policy: Frequently Asked Questions*, which is available from [EMC Online Support](#), provides more information.

Assigning users

The Avamar software includes access audit logging capability. The broad intent of this feature is to maintain a log of every action taken on vital system components/objects.

The data in this log enables enterprises that deploy Avamar to perform the following functions:

- Enforce security policies
- Follow up on security breaches or deviations from policies
- Hold appropriate users accountable for those actions

Best practices for assigning users

- Assign each Avamar administrator, operator, or user a unique login credential. Ensure that all users log in to the Avamar system by using those unique login credentials rather than the default Avamar application root and MCUser users.
- Work with EMC Customer Service to set up External Authentication so that all Avamar administrators, operators, and users can log in to the Avamar server with Active Directory, LDAP, or NIS login credentials. The *EMC Avamar Administration Guide* provides more information.

CHAPTER 11

Using Data Domain Systems

This chapter includes the following topics:

- [Network bandwidth recommendations](#) 86
- [Configuration best practices](#) 89

Network bandwidth recommendations

Before you add a Data Domain system to an Avamar configuration, ensure that the infrastructure provides adequate network bandwidth for backups and Avamar maintenance activities.

Network bandwidth in an Avamar configuration has the most impact on Avamar client backups to a Data Domain system and Avamar server maintenance activities. The process that sends Avamar client metadata to the Avamar server has less impact on the network bandwidth.

To measure the network bandwidth between the Avamar server and Data Domain system, use the `iperf` utility. The `iperf` utility is available on the Avamar server, on the Data Domain system, and from the Internet:

- On an Avamar server, the Linux operating system includes the `iperf` utility in `/usr/local/avamar/bin`.
- On a Data Domain system, the Data Domain Operating System (DD OS) includes the `iperf` utility.
- For Avamar clients, download the `iperf` utility from the Internet.

Use the iperf utility to test the network bandwidth

For the most comprehensive results, run the `iperf` utility in server mode on the Avamar server and in client mode on the Data Domain system. Then run the `iperf` utility in server mode on the Data Domain system and in client mode on the Avamar server. Run the `iperf` utility several times to verify the consistency of the results.

Procedure

1. Run the `iperf` utility in server mode on the Avamar server by typing the following command:

```
iperf -s -w 256k
```

2. Run the `iperf` utility in client mode on the Avamar server by typing the following command:

```
iperf -c iperf-server-name -w 256k
```

where *iperf-server-name* is the `iperf` server.

To view statistics for every second, add the `-i 1` option.

3. Run the `iperf` utility in server mode on the Data Domain system by typing the following command:

```
net iperf server window-size 256K
```

4. Run the `iperf` utility in client mode on the Data Domain system by typing the following command:

```
net iperf client iperf-server-name window-size 256K
```

where *iperf-server-name* is the `iperf` server.

To view statistics for every second, add the `interval 1` option.

Recommended network bandwidth

For a 1-gigabit connection to the Data Domain server, a network bandwidth of 800 Mbps/sec or greater is sufficient for client backups. A number less than 800 Mbps/sec can cause a network bottleneck and limit the throughput to a Data Domain system.

For a 10-gigabit connection to the Data Domain server, the network bandwidth of 5 Gbps/sec or greater is sufficient for client backups. A number less than 3 Gbps/sec can cause a network bottleneck on certain Data Domain systems models.

If the network bandwidth results are insufficient for the Avamar client and Data Domain system, review the information on the **Status > Stats** page in the Data Domain Enterprise Manager. This page shows network, nfs, and disk throughput. The *EMC DD OS 5.1 Administration Guide* provides more information about viewing system statistics.

Example iperf utility sessions

The following examples display output from the `iperf` utility when you use `iperf` to test bidirectional-network bandwidth between a Data Domain system and an Avamar client.

Data Domain as iperf client and Avamar client as iperf server

The following example shows the `iperf` utility output from a Data Domain as the `iperf` client and an Avamar client as the `iperf` server.

Example criteria:

- `iperf` server—Avamar client
- `iperf` server—Avamar client
- Connection—1 gigabit

```
sysadmin@datadomain1# net iperf client clidev02.lab.com window-size 256K
```

```
-----
Client connecting to clidev02.lab.com, TCP port 5001
TCP window size: 512 KByte (WARNING: requested 256 KByte)
-----
```

```
[ 3] local 192.168.0.10 port 56276 connected with 192.168.0.11
port 5001
[ 3] 0.0-10.0 sec 1.09 GBytes 938 Mbits/sec
sysadmin@datadomain1#
```

```
root@clidev02# iperf -s -w 256k
```

```
-----
Server listening on TCP port 5001
TCP window size: 264 KByte (WARNING: requested 256 KByte)
-----
```

```
[ 4] local 192.168.0.11 port 5001 connected with 192.168.0.10
port 56276
[ ID] Interval Transfer Bandwidth
[ 4] 0.0-10.0 sec 1.09 GBytes 938 Mbits/sec
```

Data Domain as iperf server and Avamar client as iperf client

The following example shows the `iperf` utility output from a Data Domain as the `iperf` server and an Avamar client as the `iperf` client.

Example criteria:

- `iperf server`—Data Domain system
- `iperf client`—Avamar client
- Connection—1 gigabit

```
sysadmin@datadomain1# net iperf server window-size 256K
-----
Server listening on TCP port 5001
TCP window size: 512 KByte (WARNING: requested 256 KByte)
-----
[ 4] local 192.168.0.10 port 5001 connected with 192.168.0.11
port 52347
[ 4] 0.0-10.0 sec 1.09 GBytes 937 Mbits/sec
root@clidev02# iperf -c datadomain1.lab.com -w 256k
-----
Client connecting to datadomain1.lab.com, TCP port 5001
TCP window size: 264 KByte (WARNING: requested 256 KByte)
-----
[ 3] local 192.168.0.11 port 52347 connected with 192.168.0.10
port 5001
[ ID] Interval Transfer Bandwidth
[ 3] 0.0-10.0 sec 1.09 GBytes 936 Mbits/sec
```

Data Domain as `iperf client` and Avamar client as `iperf server`

The following example shows the `iperf` utility output from a Data Domain as the `iperf client` and an Avamar client as the `iperf server`.

Example criteria:

- `iperf server`—Avamar client
- `iperf client`—Data Domain system
- Connection—10 gigabit

```
sysadmin@datadomain4# net iperf client clidev02.lab.com window-
size 256K
-----
Client connecting to clidev02.lab.com, TCP port 5001
TCP window size: 512 KByte (WARNING: requested 256 KByte)
-----
[ 3] local 192.168.0.12 port 37368 connected with 192.168.0.11
port 5001
[ 3] 0.0-10.0 sec 7.82 GBytes 6.71 Gbits/sec
sysadmin@datadomain4#

root@clidev02# iperf -s -w 256k
-----
Server listening on TCP port 5001
TCP window size: 264 KByte (WARNING: requested 256 KByte)
-----
[ 4] local 192.168.0.11 port 5001 connected with 192.168.0.12
port 37368
[ ID] Interval Transfer Bandwidth
[ 4] 0.0-10.0 sec 7.82 GBytes 6.71 Gbits/sec
```

Data Domain as the `iperf server` and Avamar client as `iperf client`

The following example shows the `iperf` utility output from a Data Domain as the `iperf server` and an Avamar client as the `iperf client`.

Example criteria:

- `iperfserver`—Data Domain system
- `iperf client`—Avamar client
- Connection—10 gigabit

```
sysadmin@datadomain4# net iperf server window-size 256K
-----
Server listening on TCP port 5001
TCP window size: 512 KByte (WARNING: requested 256 KByte)
-----
[ 4] local 192.168.0.12 port 5001 connected with 192.168.0.11
port 52351
[ 4] 0.0-10.0 sec 9.70 GBytes 8.33 Gbits/sec
root@clidev02# iperf -c datadomain4.lab.com -w 256k
-----
Client connecting to datadomain4.lab.com, TCP port 5001
TCP window size: 264 KByte (WARNING: requested 256 KByte)
-----
[ 3] local 192.168.0.11 port 52351 connected with 192.168.0.12
port 5001
[ ID] Interval Transfer Bandwidth
[ 3] 0.0-10.0 sec 9.70 GBytes 8.32 Gbits/sec
root@clidev02#
```

Configuration best practices

The following topics cover configuration best practices for Data Domain systems and Data Domain Archivers.

Use fully qualified domain names

Assign fully qualified domain names (FQDN) to the Data Domain system before you add it to an Avamar configuration. Adhere to the following guidelines when you configure the Data Domain system:

- Do not use the IP address of the Data Domain system. Use of an IP address instead of the hostname can limit the ability to route duplication traffic.
- Use forward and reverse DNS lookup for the Avamar server, Avamar clients, and the Data Domain system.
- Use DNS to resolve hostnames to routable IP addresses.
- Use host files to resolve hostnames to non-routable IP addresses.
- Do not create secondary hostnames for alternate or local IP interfaces.

Review the amount of files in the MTree

The DD Boost and DD OS are not designed to handle backups of large amounts of small files or backups of thousands of clients. The best use for Avamar integration with Data Domain systems is in data centers that back up a few large files.

When an MTree contains many small files, the following known issues can occur:

- Avamar checkpoint and rollback processes on a Data Domain system can fail if the Data Domain system contains more than 400,000 directories or files below the cur directory or in any checkpoint directory.

- A DD OS 4.9 conversion of an LSU to an MTree can fail if there are over 400,000 objects in the LSU.

The expiration of a two-minute time-out period causes both of these issues.

Do not modify the MTree

The DD OS protects an MTree so that only the DD Boost account has read/write privileges to the MTree. Do not modify the MTree structure by using SSH from the DD OS command line, NFS, CIFS, or the Data Domain Enterprise Manager. Modifications to the MTree by using any one of these methods can result in failed backups, restores, or Avamar server maintenance operations.

The Avamar software controls additions, modifications, and deletions to the Data Domain system by using the DD Boost library. Do not manually add, modify, or delete the contents of an MTree. Doing so can cause irreversible consistency problems between the Avamar server and Data Domain system.

Specify the maximum number of data streams

Avamar clients that support Data Domain systems as a storage device can use multiple data streams during backups and restores. For example, the Avamar Plug-in for SQL Server, the Avamar Plug-in for SharePoint VSS, and the Avamar Plug-in for Exchange VSS use one stream for each backup or restore. The Avamar Plug-in for Oracle can use 1 to 6 streams.

The number of streams you specify depends on the number of backups that you plan to run simultaneously. Adhere to the following guidelines when you add a Data Domain system to an Avamar configuration:

- Specify the total amount of streams that the specific model of the Data Domain supports if only one Avamar server uses the Data Domain system.
- Specify a subset of streams that the specific model of the Data Domain system supports if multiple Avamar servers or other third-party applications share the Data Domain system.

Evaluate storage requirements

When you configure Avamar to use a Data Domain system that third-party applications also use, carefully evaluate the amount of storage you need for Avamar data. Ensure that enough storage is available for both Avamar data and the third-party applications.

Synchronize the time

Use an Network Time Protocol (NTP) to synchronize the system time on the Avamar server and the Data Domain system.

Restore backups with the Use SQL REPLACE option

When you restore SQL Server backups from a Data Domain system, always select the **Use SQL REPLACE** option checkbox and clear the **Tail-log backup** checkbox.

The *EMC Avamar for SQL Server User Guide* provides more information about these options.

Space requirements for replication configurations

Avamar replication configurations that include Data Domain systems or Data Domain Archivers must have enough storage space to accommodate replication. After replicating

data to a Data Domain system or Data Domain Archiver, ensure that you have at least 10% of free space.

Data movement policies

To best use the storage space on a Data Domain Archiver, do not configure a data movement policy unless you thoroughly understand this feature. The following Data Domain Archiver use case provides more information.

Use case for Data Domain Archiver

A user wants to store Avamar backups that have long retention periods on a Data Domain Archiver. The user performs the following steps:

- Configures an Avamar dataset to expire in seven years.
- Creates a data movement policy on the Data Domain Archiver for the MTree that hosts the Avamar backups.
- Specifies an age threshold of six months in the data movement policy for the Avamar backup files.

The six-month age threshold setting causes Data Domain Archiver to move Avamar backups on the active tier to the archive tier. When the archive tier is full, Data Domain Archiver seals the archive tier. Sealed archive tiers are read-only.

When the Data Domain Archiver seals an archive tier that contains Avamar backups, the backups expire on the Avamar server. The Data Domain Archiver, therefore, cannot reclaim the space that was allocated for the Avamar backups.

The *EMC DD860 Archiver Administration Guide DD OS 5.1* provides more information about configuring the data movement policy.

CHAPTER 12

Using Isilon Storage Devices

This chapter includes the following topics:

• Isilon overview	94
• Configuration	94
• Backups	96
• Restores	102
• Concurrent backups and restores	103
• Performance scalability metrics	103

Isilon overview

EMC® Isilon® is a scale-out network-attached storage (NAS) system, commonly referred to as a cluster. A single Isilon storage cluster comprises three to 144 nodes. These nodes are rack-mountable servers with memory, CPU, networking, Non-Volatile Random Access Memory (NVRAM), InfiniBand® interconnects, and disks.

Each server node is built by using commodity hardware, memory, and disks. A private InfiniBand network backend connects the nodes. The Isilon OneFS® operating system powers the Isilon hardware. The OneFS operating system enables data protection, automated data balancing and migration, and the ability to add storage and performance without system downtime. In addition, you do not need to create partitions or volumes. OneFS is a fully distributed single-namespace file system that spans all nodes in the cluster. The file system is accessible by clients that connect to any node in the cluster.

To the user, all the storage from all nodes in a cluster looks like a single big volume. The amount of storage space on an Isilon system can scale up to 20 PB.

Note

$4 \text{ TB/drive} * 36 \text{ drives/node} * 144 \text{ nodes} = 20.7 \text{ PB}$

Isilon storage systems store unstructured data in high performance computing environments. Some of the industries that use Isilon storage systems include the following types:

- Financial services
- Internet and hosting services
- Business intelligence
- Engineering
- Manufacturing
- Media and entertainment
- Bioinformatics
- Scientific research

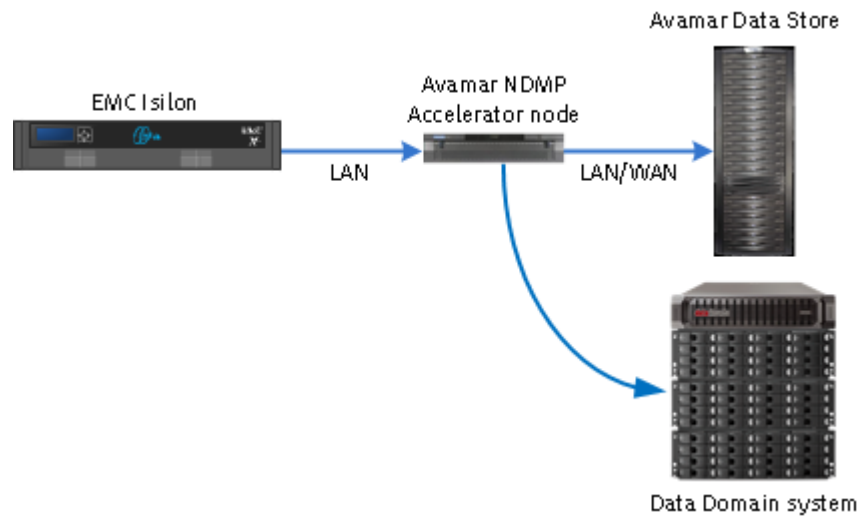
Types of unstructured data include files, such as video, audio, image, MS Word, Excel worksheets, and so forth. Types of structured data include database objects, such as tables, indexes, views, and so forth.

Specific details about all Isilon platforms are available from [EMC Online Support](#).

Configuration

This topic describes how to configure an Isilon storage cluster with one or more Avamar NDMP Accelerator nodes.

The following figure shows a simple configuration that includes an Isilon storage cluster, one Avamar NDMP Accelerator node, one Avamar Data Store and one Data Domain system.

Figure 6 Isilon configuration with Avamar and Data Domain

Accelerator nodes are available with up to 36 GB of RAM for regular and high-density file systems. The following table lists configuration requirements for the environment that includes a 300 TB Isilon cluster, accelerator node, an Avamar Data Store system, and a Data Domain system.

Table 18 Requirements for Isilon NAS device and Avamar configuration

Resource	Requirements
Connection between the Isilon cluster and the Avamar NDMP Accelerator nodes	LAN
Connection between the Avamar NDMP Accelerator node and the Avamar Data Store	LAN or WAN
Network interface	1000BaseT between NAS device, Avamar NDMP Accelerator nodes, Avamar Data Store, and Data Domain.
Number of Avamar NDMP Accelerator nodes	<ul style="list-style-type: none"> Minimum of 1 accelerator nodes per every 1x4 Avamar Data Store configuration. Maximum of 4 accelerator nodes per every 1x16 Avamar Data Store configuration.

The *EMC Avamar NDMP Accelerator for EMC NAS Systems User Guide* provides more information about configuring the Isilon storage cluster for use with Avamar NDMP Accelerator.

Guidelines for configuring the Avamar environment

- Be aware that sharing the Avamar Data Store with other clients can have negative effects. An Avamar Data Store that you use for client backups other than Isilon backups, affects ingestion rate and capacity:
 - Ingestion rate

If you run other client backups at the same time as an Isilon backup, then the ability to ingest Isilon data is affected. The issue is more pronounced for level 0 backups than for level 1 backups.

- Capacity

If you store client data on the Avamar Data Store other than Isilon data, the available capacity is limited for Isilon data.

- Limit the number of accelerators to 4:

- The number of Avamar nodes impacts the number of accelerator nodes that you can configure.
- Use of more than 4 streams per accelerator provides minimal incremental improvement.

[Performance scalability metrics on page 103](#) provide more information.

Backups

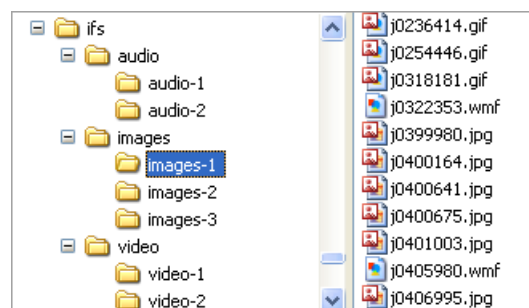
You can perform both on-demand and scheduled backups by using the Avamar NDMP Accelerator plug-in with Avamar Administrator.

Note

To perform on-demand backups, you must create a browse user account on the Isilon system. The *EMC Avamar NDMP Accelerator for EMC NAS Systems User Guide* provides more information about creating this account.

The root directory on an Isilon storage cluster is named `/ifs`. A tier of directories below `/ifs` are known as the top-level directories (TLDs). The following figure shows the `/ifs` with three TLDs: audio, images, and video.

Figure 7 Isilon root directory and three TLDs



To back up data from the Isilon file system, you select one or more TLDs or subdirectories in a TLD to back up specific directories. Be cautious when you select subdirectories in a TLD for backup. A backup of a subdirectory does not include files that are located above the subdirectory.

Avamar Administrator enables you to browse the file system below the volume level (`/ifs`). Avamar Administrator, however, cannot display all entries in directories with more than 1000 entries.

The *EMC Avamar NDMP Accelerator for EMC NAS Systems User Guide* provides more information about backing up data.

Best practices for backing up Isilon

- Avoid backing up `/ifs` in its entirety as one backup. A backup of `/ifs`, depending on its size, can take a long time to complete.
- To back up files in `/ifs`, you must exclude all TLDs and subdirectories. [Exclude lists for Isilon backups on page 97](#) provides more information about creating exclude lists.
- Do not back up a single file. You can, however, put a single file in a subdirectory, and then back up the subdirectory.
- Do not include hardlinks across separately backed-up directories.
The backup of a file, which has hardlinks across multiple directories, includes multiple copies of the file. When you restore this backup, you restore multiple copies of the same file. The restore does not handle hardlinks across backups.
- Be careful when renaming subdirectories or moving TLDs. If you move or rename a subdirectory or TLD, the next incremental backup includes the directories that you moved and all descendents.
For example, if you rename a directory from `Videos/Projects` to `Videos/Tasks`, the next incremental backup you perform will resort to a full backup.
- Directories must be organized to distribute cumulative size evenly whenever possible.
- To provide application consistent recovery, the application data must reside in a single directory tree.

Exclude lists for Isilon backups

Use an exclude list to specify one or more directories from an Isilon backup.

You can specify an exclude list for an on-demand backup, in a dataset for a scheduled backup, or with the `avndmp` command from the command line. The directory path must conform to the following format:

```
root-dir:subdir[,subdir,...]
```

where:

- *root-dir* is the root directory. Separate the root directory from subdirectories by a colon (:).
- *subdir* is a subdirectory. Separate multiple subdirectories by commas (,).

The length of a directory path must not exceed 1024 characters.

The following example shows a properly-formed directory path that excludes two directories, `/ifs/backup_dir2/relative_path_2_1` and `/ifs/backup_dir2/relative_path_2_2`:

```
/ifs/backup_dir2:/relative_path_2_1,/relative_path_2_2
```

An exclude list you create in a dataset for an incremental backup must be the same as an exclude list you create in a dataset for a full backup.

When full and incremental exclude lists differ, the course of action an incremental backup takes depends on the type of incremental backup you specify.

- If you specify the **Prefer incremental, but do a Full if required** option, a full backup runs.

- If you specify the **Force an incremental (level 1) dump** option, the incremental backup fails.

Note

To ensure the exclude lists for full and incremental backups are exactly the same, use a copy and paste strategy.

Setting up an exclude list for an on-demand backup

To exclude directories from an on-demand backup, you can specify the `--[avndmp]isilon-exclude` flag and one or more directories to exclude by using the **More** button in the **Backup Command Line Options** dialog box.

Procedure

1. From the **Backup Command Line Options** dialog box, click **More**.
The dialog box expands to display the **Enter Attribute** and **Enter Attribute Value** fields.
2. In the **Enter Attribute** field, type the following flag:

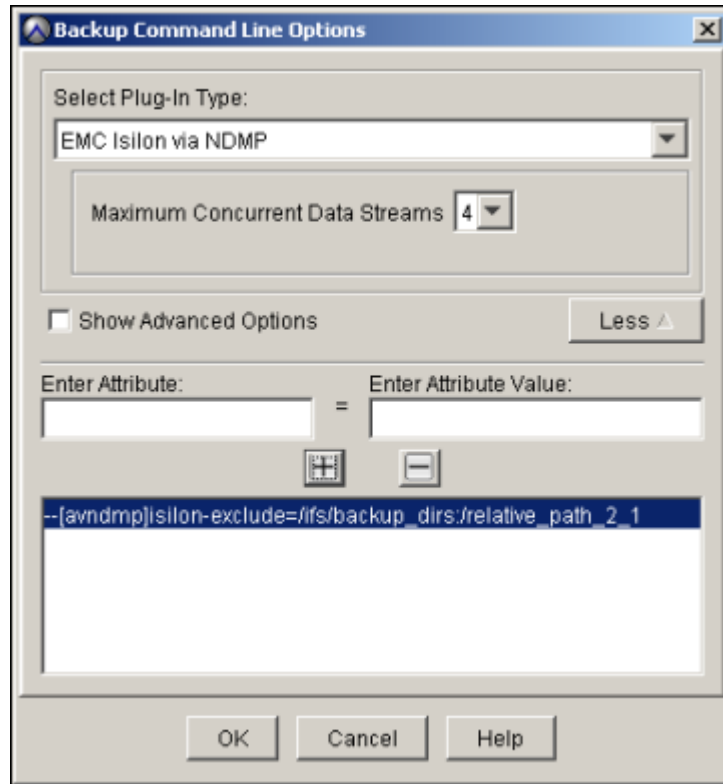
```
--[avndmp]isilon-exclude
```

3. In the **Enter Attribute Value** field, type the exclude list. For example to exclude the `/ifs/backup_dir2/relative_path_2_1` from the backup, you would type the directory path by using the following format:

```
/ifs/backup_dir2:/relative_path_2_1
```

4. Click **+**.

The `--[avndmp]isilon-exclude` flag and value appear in the area below the **Enter Attribute** and **Enter Attribute Value** fields. The following figure shows the **Backup Command Line Options** dialog box after you add the `--[avndmp]isilon-exclude` flag and value.



Setting up an exclude list for a scheduled backup

To exclude directories from a scheduled backup, you can specify the `--[avndmp]isilon-exclude` flag and one or more directories to exclude by using the **Options** tab in the **New Dataset** dialog box.

Procedure

1. From the **New Dataset** dialog box, select the **Options** tab.
The dialog box expands to display the **Enter Attribute** and **Enter Attribute Value** fields.
2. In the **Enter Attribute** field, type the following flag:

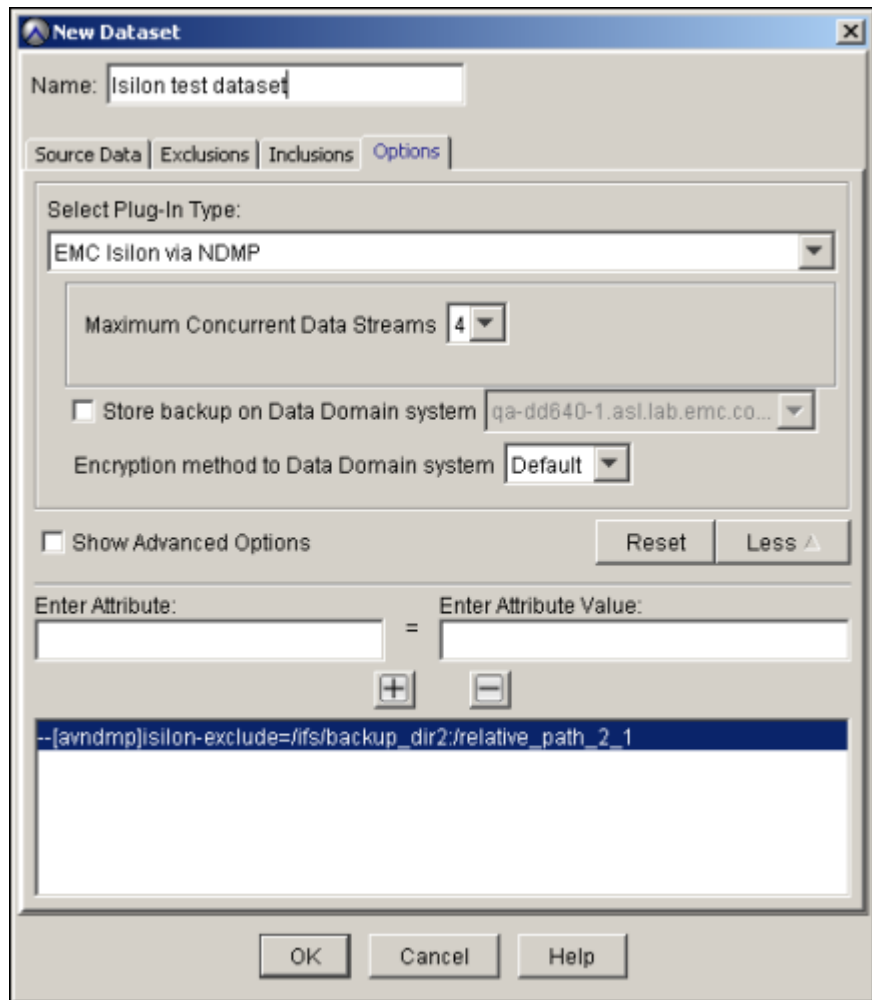
```
--[avndmp]isilon-exclude
```

3. In the **Enter Attribute Value** field, type the exclude list. For example to exclude the `/ifs/backup_dir2/relative_path_2_1` from the backup, you would type the directory by using the following format:

```
/ifs/backup_dir2:/relative_path_2_1
```

4. Click **+**.

The option and value appear in the area below the **Enter Attribute** and **Enter Attribute Value** fields. The following figure shows the **New Dataset** dialog box after you add the `--[avndmp]isilon-exclude` option and value.



5. Click **OK**.

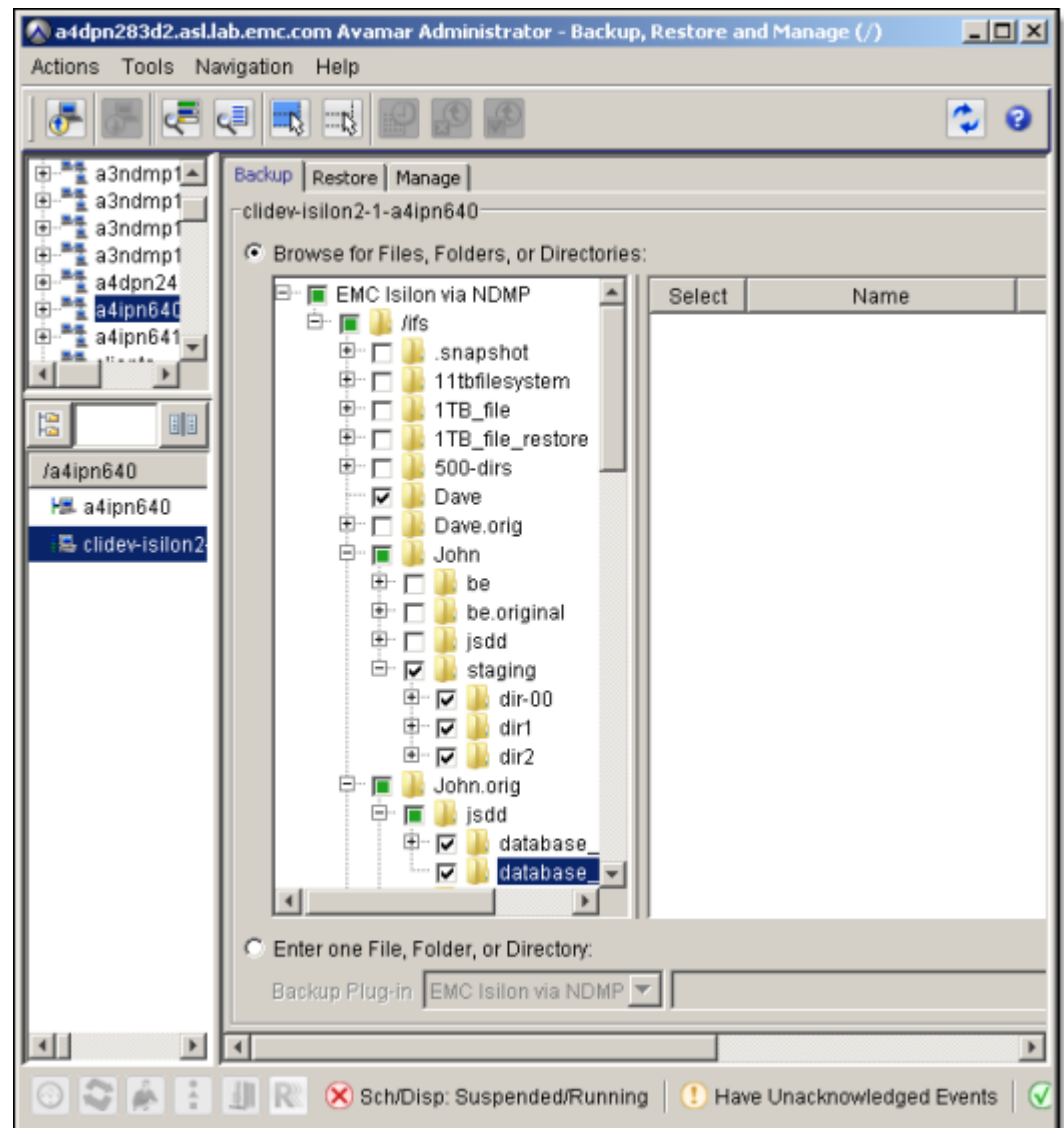
Note

The Avamar NDMP Accelerator plug-in does not support the creation of exclude or include lists from the **Exclusions** tab or **Inclusions** tab of the **New Dataset** dialog box.

Multi-streaming backups

You can back up multiple directories in parallel by specifying the number of data streams for a backup. The default number of streams is 4. The recommended number of streams is 8. You can reduce the number of streams if there is too much load on the server.

Multi-streaming backups can comprise TLDs, TLDs and subdirectories, or just subdirectories. The following figure shows the **Backup** tab after the user selects two TLDs and three subdirectories for a backup.

Figure 8 Backup selection comprising TLDs and subdirectories

Avamar Administrator displays a checkbox next to items available for backup.

- A checkbox with a check mark indicates that the item is marked for backup.
- A checkbox that is filled with the color green indicates that an item within the directory is marked for backup.
- An empty checkbox indicates that the item is not marked for backup.

You can specify multiple streams for an on-demand backup or in a dataset for a scheduled backup by selecting the number of streams from the **Maximum Concurrent Data Streams** list. When you run a multi-stream backup, Avamar runs a separate backup for each stream. For example, if you select 8 items for a backup and select 8 streams, Avamar runs 8 backups in parallel. If you select 10 items for a backup and select 8 streams, Avamar runs 8 backups in parallel. The 2 remaining backups start as soon as a backup stream becomes available.

You must define a multi-streaming backup at the TLD.

Replication target backups

You can back up a replication target for an Isilon storage cluster by using the Avamar NDMP Accelerator plug-in.

Backups after renaming a TLD

After you rename a TLD, you must perform a full backup. Before you perform the full backup, add the new TLD name to the list of TLDs to back up and remove the old name from the backup list.

Backups after moving a TLD inside an existing TLD

After you move a previously backed up TLD inside an existing TLD, the next incremental backup for the existing TLD includes the full contents of the moved TLD in the incremental data stream. You must remove the name of the “moved” TLD from the backup list.

Restores

You can restore data to an Isilon storage cluster by using the Avamar NDMP Accelerator plug-in with Avamar Administrator. The **Backup, Restore and Manage** window in Avamar Administrator enables you to find an Isilon backup by specifying a date or by specifying a file or a directory path. The second method requires that you know the name of the file or directory.

The *EMC Avamar NDMP Accelerator for EMC NAS Systems User Guide* provides more information about restoring data.

Best practice for restoring Isilon

Do not restore the `/ifs` root directory in its entirety. Avoid performing a restore of `/ifs` in its entirety because the `/ifs/.ifsvvar` directory contains system configuration files. A restore to the `/ifs/.ifsvvar` directory can overwrite the existing system configuration and cause cluster problems.

Multi-streaming restores

The Avamar NDMP Accelerator plug-in enables you to specify 1 to 8 multiple data streams for a restore operation. You specify the number of streams for the restore by selecting a value from the **Maximum Concurrent Data Streams** list in the **Restore Command Line Options** dialog box.

Cross-platform restores

The Avamar NDMP Accelerator plug-in enables you to perform cross-platform restores. A cross-platform restore restores an Isilon backup directly to CIFS mounts on a Windows system or NFS mounts on a Linux system. A cross-platform restore does not restore folder and file ACLs, or Windows alternate data streams.

A cross-platform restore handles sparse files by only backing up blocks with data and skipping the sparse blocks. The restore uses the sparse map from the backup to properly recover the data blocks.

You perform a cross-platform restore by selecting the **Restore everything to a different location** option in the **Restore Options** dialog box and specifying the target client for the

restore. The *EMC Avamar NDMP Accelerator for EMC NAS Systems User Guide* provides more information about restoring Isilon data to a different client.

Concurrent backups and restores

For backups and restores, Isilon supports up to 64 concurrent NDMP sessions per node. EMC recommends 8 concurrent NDMP backups or restores.

Each backup session may use up to 512 MB of memory on the Isilon system. Too many concurrent NDMP backup sessions may run out of memory and slow down the backups.

Performance scalability metrics

Optimal backup and restore performance depends on the hardware configuration, the size of the TLDs, and the number of backup or restore data streams.

The initial backup of an Isilon cluster at the TLD level results in a full (level 0) backup. All subsequent incremental (level 1) backups save all data that has changed since the last backup. Those level 1 backups are synthesized into a full image for recovery purposes.

The test results that follow can help you determine the best backup and restore strategy to obtain optimal performance.

Isilon backups to and restores from Avamar Data Store

The following table lists configuration information for the test environment that was used to obtain performance metrics for backups to and restores from an Avamar Data Store.

Table 19 Configuration for testing Isilon backups to and restores from an Avamar Data Store

Hardware	Configuration details
Isilon system	<ul style="list-style-type: none"> OneFS v7.0.1.2 3 cluster nodes (X200 2U Single 24 GB) 55.2 TB total storage
Avamar NDMP Accelerator	<ul style="list-style-type: none"> 3 x Gen 4S 32 GB memory 2 x CPU Avamar NDMP Accelerator version: 7.0.100
Avamar Data Store	<ul style="list-style-type: none"> Gen 4S 7.8 TB Avamar server version: 7.0.0

The following list describes the data characteristics for these performance tests:

- 4 TB dataset size with 100% unique, uncompressible data
- 5 MB average file size
- Data distributed across 32 directories
- Each directory size = 128 GB

- Approximately 800,000 files and 60,000 directories

The following tables extrapolate performance results for backups and restores of a 100 TB Isilon cluster. These results are from backup and restore operations performed in a lab environment. Actual backup and restore performance results in your environment can vary.

The following table shows backup performance results for level 0 backups.

Table 20 Performance results for level 0 backups to an Avamar Data Store

Avamar Data Store type	Number of accelerators	Number of streams	Throughput	Backup time
1 x 4	1	4	320 GB/hr	13 days
1 x 8	1	4	345 GB/hr	12 days
1 x 8	2	8	574 GB/hr	7.2 days
1 x 8	3	8	577 GB/hr	7.2 days

The following table shows backup performance results for level 1 backups. Backup data comprised 10% changed and 1% new.

Table 21 Performance results for level 1 backups to an Avamar Data Store

Avamar Data Store type	Number of accelerators	Number of streams	Effective throughput	Actual throughput	Backup time
1 x 8	1	8	371 GB/hr	39.01 GB/hr	25 hours
1 x 8	2	8	743 GB/hr	77.64 GB/hr	13 hours
1 x 8	3	8	1069 GB/hr	116.2 GB/hr	9.3 hours

The following table shows restore performance.

Table 22 Performance results for restore operations from an Avamar Data Store

Number of streams	Number of accelerators	Throughput
8	1	261 GB/hr
8	2	573 GB/hr
8		779 GB/hr

Summary of performance testing

You can connect multiple accelerator nodes to an Isilon cluster.

Optimal accelerator node configuration:

- Level 0 backups with 2 accelerator nodes obtain approximately 570 GB/hr throughput. This throughput is obtained by using Avamar Data Store 1x8 or larger configurations.

- Level 1 backups with 4 accelerator nodes per Isilon cluster obtain the following throughput:
 - Effective throughput of 1.2 to 1.4 TBs/hr (scan plus data transfer)
 - Actual throughput of 120 to 140 GBs/hr (new data transfer)

Dense file systems (numerous files) take significantly longer to back up than file system that have less files.

