



EMC[®] NetWorker[®]
Release 7.6 Service Pack 1

Performance Optimization Planning Guide

P/N 300-011-323

REV A01

EMC Corporation

Corporate Headquarters:
Hopkinton, MA 01748-9103

1-508-435-1000

www.EMC.com

Copyright © 1990-2010 EMC Corporation. All rights reserved.

Published September, 2010

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.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.

Preface

Chapter 1

Overview

Organization	10
NetWorker data flow	11

Chapter 2

Size the NetWorker Environment

Expectations	14
Determine backup environment performance expectations	14
Determine required backup expectations	14
System components	16
System	16
Storage	19
Network	20
Target device	20
The component 70 percent rule	21
Components of a NetWorker environment	22
Datazone	22
NetWorker Management Console	22
NetWorker server	23
NetWorker storage node	24
NetWorker client	25
NetWorker databases	25
Optional NetWorker Application Modules	26
Virtual environments	26
NetWorker deduplication nodes	26
Recovery performance factors	27
Connectivity and bottlenecks	28
NetWorker database bottlenecks	32

Chapter 3

Tune Settings

Optimize NetWorker parallelism	36
Server parallelism	36
Client parallelism	36
Group parallelism	36
Multiplexing	37
Device performance tuning methods	38

Input/output transfer rate.....	38
Built-in compression.....	38
Drive streaming.....	38
Device load balancing	38
Network devices.....	39
DataDomain.....	39
AFTD device target and max sessions.....	40
Number of virtual device drives versus physical device drives.....	40
Network optimization	42
Advanced configuration optimization.....	42
Operating system TCP stack optimization.....	42
Advanced tuning	43
Network latency	43
Ethernet duplexing	45
Firewalls	45
Jumbo frames.....	45
Congestion notification	45
TCP buffers	46
NetWorker socket buffer size.....	47
IRQ balancing and CPU affinity	47
TCP offloading	48
Name resolution.....	48
Storage optimization.....	49
NetWorker server and storage node disk write latency	49

Chapter 4 Test Performance

Determine symptoms	52
Monitor performance.....	53
Determine bottlenecks by using a generic FTP test.....	54
Test the performance of the setup by using dd.....	55
Test disk performance by using bigasm and uasm	56
The bigasm directive	56
The uasm directive.....	56

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

If a product does not function properly or does not function as described in this document, please contact your EMC representative.

Audience This document is part of the NetWorker documentation set, and is intended for use by system administrators to identify the different hardware and software components that make up the NetWorker datazone. It discusses the component's impact on storage management tasks, and provides general guidelines for locating problems and solutions.

NetWorker product documentation

This section describes the additional documentation and information products that are available with NetWorker.

EMC NetWorker Release 7.6 Service Pack 1 Installation Guide

Provides instructions for installing or updating the NetWorker software for clients, console and server on all supported platforms.

EMC NetWorker Release 7.6 Service Pack 1 Cluster Installation Guide

Contains information related to installation of the NetWorker software on cluster server and clients.

EMC NetWorker Release 7.6 Service Pack 1 Administration Guide

Describes how configure and maintain the NetWorker software.

EMC NetWorker Release 7.6 Service Pack 1 Release Notes

Contain information on new features and changes, fixed problems, known limitations, environment and system requirements for the latest NetWorker software release.

NetWorker Data Domain Deduplication Devices Integration Guide

Provides planning and configuration information on the use of Data Domain devices for data deduplication backup and storage in a NetWorker environment.

EMC NetWorker Licensing Guide

Provides information about licensing NetWorker products and features.

NetWorker License Manager 9th Edition Installation and Administration Guide

Provides installation, set up, and configuration information for the NetWorker License Manager product.

NetWorker 7.6 Service Pack 1 Error Message Guide

Provides information on common NetWorker error messages.

NetWorker 7.6 Service Pack 1 Command Reference Guide

Provides reference information for NetWorker commands and options.

NetWorker Management Console Online Help

Describes the day-to-day administration tasks performed in the NetWorker Management Console and the NetWorker Administration window. To view Help, click Help in the main menu.

NetWorker User Online Help

The NetWorker User program is the Windows client interface. Describes how to use the NetWorker User program which is the Windows client interface connect to a NetWorker server to back up, recover, archive, and retrieve files over a network.

NetWorker related documentation

For more information about NetWorker software, refer to this documentation:

EMC Information Protection Software Compatibility Guide

A list of supported client, server, and storage node operating systems for the following software products: AlphaStor, ArchiveXtender, DiskXtender for Unix/Linux, DiskXtender for Windows, Backup Advisor, AutoStart, AutoStart SE, RepliStor, NetWorker, and NetWorker Modules and Options.

E-lab Issue Tracker

Issue Tracker offers up-to-date status and information on NetWorker known limitations and fixed bugs that could impact your operations. E-Lab Issue Tracker Query allows you to find issues in the Issue Tracker database by matching issue number, product feature, host operating system, fixed version, or other fields.

NetWorker Procedure Generator

The NetWorker Procedure Generator (NPG) is a stand-alone Windows application used to generate precise user driven steps for high demand tasks carried out by customers, support and the field. With the NPG, each procedure is tailored and generated based on user-selectable prompts. This generated procedure gathers the most critical parts of NetWorker product guides and combines experts' advice into a single document with a standardized format.

Note: To access the E-lab Issue Tracker or the NetWorker Procedure Generator, go to <http://www.Powerlink.emc.com>. You must have a service agreement to use this site.

Technical Notes and White Papers

Provides an in-depth technical perspective of a product or products as applied to critical business issues or requirements. Technical Notes and White paper types include technology and business considerations, applied technologies, detailed reviews, and best practices planning.

Conventions used in this document

EMC uses the following conventions for special notices.

Note: A note presents information that is important, but not hazard-related.

**CAUTION**

A caution contains information essential to avoid data loss or damage to the system or equipment.

**IMPORTANT**

An important notice contains information essential to operation of the software.

Typographical conventions

EMC uses the following type style conventions in this document:

Normal	Used in running (nonprocedural) text for: <ul style="list-style-type: none"> Names of interface elements (such as names of windows, dialog boxes, buttons, fields, and menus) Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, functions, utilities URLs, pathnames, filenames, directory names, computer names, filenames, links, groups, service keys, file systems, notifications
Bold	Used in running (nonprocedural) text for: <ul style="list-style-type: none"> Names of commands, daemons, options, programs, processes, services, applications, utilities, kernels, notifications, system calls, man pages
	Used in procedures for: <ul style="list-style-type: none"> Names of interface elements (such as names of windows, dialog boxes, buttons, fields, and menus) What user specifically selects, clicks, presses, or types
<i>Italic</i>	Used in all text (including procedures) for: <ul style="list-style-type: none"> Full titles of publications referenced in text Emphasis (for example a new term) Variables
<code>Courier</code>	Used for: <ul style="list-style-type: none"> System output, such as an error message or script URLs, complete paths, filenames, prompts, and syntax when shown outside of running text
<code>Courier bold</code>	Used for: <ul style="list-style-type: none"> Specific user input (such as commands)
<i><code>Courier italic</code></i>	Used in procedures for: <ul style="list-style-type: none"> Variables on command line User input variables
< >	Angle brackets enclose parameter or variable values supplied by the user
[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections - the bar means "or"
{ }	Braces indicate content that you must specify (that is, x or y or z)
...	Ellipses indicate nonessential information omitted from the example

Where to get help

EMC support, product, and licensing information can be obtained as follows.

Product information — For documentation, release notes, software updates, or for information about EMC products, licensing, and service, go to the EMC Powerlink website (registration required) at: <http://Powerlink.EMC.com>

Technical support — For technical support, go to EMC Customer Service on Powerlink. To open a service request through Powerlink, you must have a valid support agreement. Please contact your EMC sales representative for details about obtaining a valid support agreement or to answer any questions about your account.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Please send your opinion of this document to:

BSGdocumentation@emc.com

If you have issues, comments, or questions about specific information or procedures, include the title and, if available, the part number, the revision (for example, A01), the page numbers, and any other details that will help us locate the subject you are addressing.

The NetWorker software is a network storage management application that is optimized for high-speed backup and recovery operations of large amounts of complex data across an entire datazone. This guide addresses non-disruptive performance tuning options. Although some physical devices may not meet the expected performance, it is understood that when a physical component is replaced with a better performing device, another component ends up as a bottle neck. This manual attempts to address NetWorker performance tuning with minimal disruptions to the existing environment. It attempts to fine-tune feature functions to achieve better performance with the same set of hardware, and to assist administrators to:

- ◆ Understand data transfer fundamentals
- ◆ Determine requirements
- ◆ Identify bottlenecks
- ◆ Optimize and tune NetWorker performance.

This chapter includes these sections:

- ◆ [Organization](#) 10
- ◆ [NetWorker data flow](#) 11

Organization

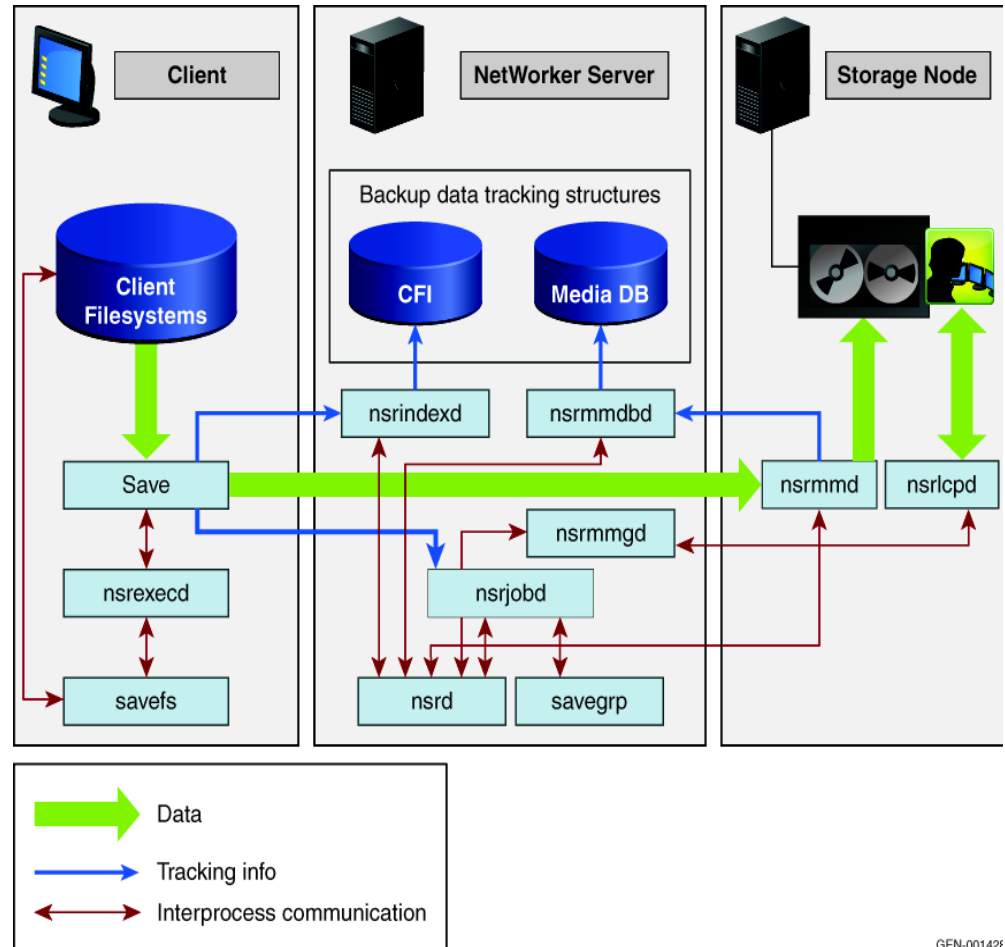
This guide is organized into the following chapters:

- ◆ [Chapter 2, “Size the NetWorker Environment,”](#) provides details on how to determine requirements.
- ◆ [Chapter 3, “Tune Settings,”](#) provides details on how to tune the backup environment to optimize backup and restore performance.
- ◆ [Chapter 4, “Test Performance,”](#) provides details on how to test and understand bottlenecks by using available tools.

NetWorker data flow


Figure 1 on page 11 and Figure 2 on page 12 illustrate the backup and recover data flow for components in an EMC® NetWorker datazone.

Note: Figure 1 and Figure 2 are simplified diagrams, and not all interprocess communication is shown. There are many other possible backup and recover data flow configurations.



GEN-001428

Figure 1 NetWorker backup data flow



GEN-001429

Figure 2 NetWorker recover data flow

This chapter describes how to best determine backup and system requirements. The first step is to understand the environment. Performance issues can often be attributed to hardware or environmental issues. An understanding of the entire backup data flow is important to determine the optimal performance that can be expected from the NetWorker software.

This chapter includes the following topics:

- ◆ Expectations 14
- ◆ System components 16
- ◆ Components of a NetWorker environment 22
- ◆ Recovery performance factors 27
- ◆ Connectivity and bottlenecks 28

Expectations

This section describes backup environment performance expectations and required backup configurations.

Determine backup environment performance expectations

Sizing considerations for the backup environment are listed here:

- ◆ Review the network and storage infrastructure information before setting performance expectations for your backup environment including the NetWorker server, storage nodes, and clients.
- ◆ Review and set the Recovery Time Objective (RTO) for each client.
- ◆ Determine the backup window for each NetWorker client.
- ◆ List the amount of data to be backed up for each client during full and incremental backups.
- ◆ Determine the data growth rate for each client.
- ◆ Determine client browse and retention policy requirements.

It is difficult to precisely list performance expectations, while keeping in mind the environment and the devices used. It is good to know the bottlenecks in the setup and to set expectations appropriately.

Some suggestions to help identify bottlenecks and define expectations are:

- ◆ Create a diagram
 - ◆ List all system, storage, network, and target device components
 - ◆ List data paths
 - ◆ Mark down the bottleneck component in the data path of each client
- [“Connectivity and bottlenecks” on page 28](#) provides examples of possible bottlenecks in the NetWorker environment.

It is very important to know how much down time is possible for each NetWorker client. This dictates the RTO. Review and document the RTO for each NetWorker client.

To determine the backup window for each client:

1. Verify the available backup window for each NetWorker client.
2. List the amount of data that must be backed up from the clients for full or incremental backups.
3. List the average daily/weekly/monthly data growth on each NetWorker client.

Determine required backup expectations

Methods to determine the required backup configuration expectations for the environment are listed here:

- ◆ Verify the existing backup policies and ensure that the policies will meet the RTO for each client.
- ◆ Estimate backup window for each NetWorker client based on the information collected.

- ◆ Determine the organization of the separate NetWorker client groups based on these parameters:
 - Backup window
 - Business criticality
 - Physical location
 - Retention policy
- ◆ Ensure that RTO can be met with the backup created for each client.

The shorter the acceptable downtime, the more expensive backups are. It may not be possible to construct a backup image from a full backup and multiple incremental backups if the acceptable down time is very short. Full backups might be required more frequently which results in a longer backup window. This also increases network bandwidth requirements.

System components

Every backup environment has a bottleneck. It may be a fast bottleneck, but the bottleneck will determine the maximum throughput obtainable in the system. Backup and restore operations are only as fast as the slowest component in the backup chain.

Performance issues are often attributed to hardware devices in the datazone. This guide assumes that hardware devices are correctly installed and configured.

This section discusses how to determine requirements. For example:

- ◆ How much data must move?
- ◆ What is the backup window?
- ◆ How many drives are required?
- ◆ How many CPUs are required?

Devices on backup networks can be grouped into four component types. These are based on how and where devices are used. In a typical backup network, the following four components are present:

- ◆ System
- ◆ Storage
- ◆ Network
- ◆ Target device

System

The components that impact performance in system configurations are listed here:

- ◆ CPU
- ◆ Memory
- ◆ System bus (this determines the maximum available I/O bandwidth)

CPU requirements

Determine the optimal number of CPUs required, if 5 MHz is required to move 1 MB of data from a source device to a target device. For example, a NetWorker server, or storage node backing up to a local tape drive at a rate of 100 MB per second, requires 1 GHz of CPU power:

- ◆ 500 MHz is required to move data from the network to a NetWorker server or storage node.
- ◆ 500 MHz is required to move data from the NetWorker server or storage node to the backup target device.

Note: 1 GHz on one type of CPU does not directly compare to a 1 GHz of CPU from a different vendor.

The CPU load of a system is impacted by many additional factors. For example:

- ◆ High CPU load is not necessarily a direct result of insufficient CPU power, but can be a side effect of the configuration of the other system components.
- ◆ Drivers:

Be sure to investigate drivers from different vendors as performance varies. Drivers on the same operating system achieve the same throughput with a significant difference in the amount of CPU used.

- ◆ Disk drive performance:
 - On a backup server with 400 or more clients in /nsr, a heavily used disk drive often results in CPU use of more than 60 percent. The same backup server in /nsr on a disk array with low utilization, results in CPU use of less than 15 percent.
 - On UNIX, and Windows if a lot of CPU time is spent in privileged mode or if a percentage of CPU load is higher in system time than user time, it often indicates that the NetWorker processes are waiting for I/O completion. If the NetWorker processes are waiting for I/O, the bottleneck is not the CPU, but the storage used to host NetWorker server.
 - On Windows, if a lot of time is spent on Deferred Procedure Calls it often indicates a problem with device drivers.
- ◆ Monitor CPU use according to the following classifications:
 - User mode
 - System mode
- ◆ Hardware component interrupts cause high system CPU use resulting poor performance. If the number of device interrupts exceed 10,000 per second, check the device.

Memory requirements

[Table 1 on page 17](#) lists the minimum memory requirements for the NetWorker server. This ensures that memory is not a bottleneck.

Table 1

Minimum required memory for the NetWorker server

Number of clients	Minimum required memory
Less than 50	4 GB
51–150	8 GB
More than 150	16 GB

Monitor the pagefile or swap use

Memory paging should not occur on a dedicated backup server as it will have a negative impact on performance in the backup environment.

System bus requirements

Although HBA/NIC placement are critical, the internal bus is probably the most important component of the operating system. The internal bus provides communication between internal computer components, such as CPU, memory, disk, and network.

Bus performance criteria:

- ◆ Type of bus
- ◆ Data width
- ◆ Clock rate
- ◆ Motherboard

System bus considerations:

- ◆ A faster bus does not guarantee faster performance
- ◆ Higher end systems have multiple buses to enhance performance
- ◆ The bus is often the main bottleneck in a system

System bus recommendations

It is recommended to use PCIeExpress for both servers and storage nodes to reduce the chance for I/O bottlenecks.

Note: Avoid using old bus types or high speed components optimized for old bus type as they generate too many interrupts causing CPU spikes during data transfers.

PCI-X and PCIeExpress considerations:

- ◆ PCI-X is a half-duplex bi-directional 64-bit parallel bus.
- ◆ PCI-X bus speed may be limited to the slowest device on the bus, be careful with card placement.
- ◆ PCIeExpress is full-duplex bi-directional serial bus using 8/10 encoding.
- ◆ PCIeExpress bus speed may be determined per each device.
- ◆ Do not connect a fast HBA/NIC to a slow bus, always consider bus requirements. Silent packet drops can occur on a PCI-X 1.0 10GbE NIC, and bus requirements cannot be met.
- ◆ Hardware that connects fast storage to a slower HBA/NIC will slow overall performance.

“The component 70 percent rule” on page 21 provides details on the ideal component performance levels.

Note: The aggregate number of bus adapters should *not* exceed bus specifications.

Bus speed requirements

Bus speed requirements are listed below:

- ◆ 4 Gb Fibre Channel requires 425 MB/s
- ◆ 8 Gb Fibre Channel requires 850 MB/s
- ◆ 10 GB Fibre Channell requires 1,250 MB/s

Bus specifications

Bus specifications are listed in [Table 2 on page 18](#).

Table 2 Bus specifications (page 1 of 2)

Bus type	MHz	MB/second
PCI 32-bit	33	133
PCI 64-bit	33	266
PCI 32-bit	66	266
PCI 64-bit	66	533
PCI 64-bit	100	800
PCI-X 1.0	133	1,067
PCI-X 2.0	266	2,134

Table 2 Bus specifications (page 2 of 2)

Bus type	MHz	MB/second
PCI-X 2.0	533	4,268
PCleXpress 1.0 x 1		250
PCleXpress 1.0 x 2		500
PCleXpress 1.0 x 4		1,000
PCleXpress 1.0 x 8		2,000
PCleXpress 1.0 x 16		4,000
PCleXpress 1.0 x 32		8,000
PCleXpress 2.0 x 8		4,000
PCleXpress 2.0 x 16		8,000
PCleXpress 2.0 x 32		16,000

Storage

The components that impact performance of storage configurations are listed here:

- ◆ Storage connectivity:
 - Local versus SAN attached versus NAS attached
 - Use of storage snapshots
 - The snapshot technology used determines the read performance
- ◆ Storage replication:
 - Some replication technologies add significant latency to write access slows down storage access.
- ◆ Storage type:
 - Serial ATA (SATA) computer bus is a storage-interface for connecting host bus adapters to storage devices such as hard disk drives and optical drives.
 - Fibre Channel (FC) is a gigabit-speed network technology primarily used for storage networking.
 - Flash is a non-volatile computer storage used for general storage and the transfer of data between computers and other digital products.
- ◆ I/O transfer rate of storage:
 - I/O transfer rate of storage is influenced by different RAID levels, where the best RAID level for the backup server is RAID1 or RAID5. Backup to disk should use RAID3.
- ◆ Scheduled I/O:
 - If the target system is scheduled to perform I/O intensive tasks at a specific time, schedule backups to run at a different time.

- ◆ I/O data:
 - Raw data access offers the highest level of performance, but does not logically sort saved data for future access.
 - File systems with a large number of files have degraded performance due to additional processing required by the file system.
- ◆ Compression:

If data is compressed on the disk, the operating system or an application, the data is decompressed before a backup. The CPU requires time to re-compress the files, and disk speed is negatively impacted.

Network

The components that impact network configuration performance are listed here:

- ◆ IP network

A computer network made of devices that support the Internet Protocol to determine the source and destination of network communication.
- ◆ Storage network

The system on which physical storage, such as tape, disk, or file system resides.
- ◆ Network speed

The speed at which data travels over the network.
- ◆ Network bandwidth

The maximum throughput of a computer network.
- ◆ Network path

The communication path used for data transfer in a network.
- ◆ Network concurrent load

The point at which data is placed in a network to ultimately maximize bandwidth.
- ◆ Network latency

The measure of the time delay for data traveling between source and target devices in a network.

Target device

The components that impact performance in target device configurations are listed here:

- ◆ Storage type:
 - Raw disk versus Disk Appliance:
 - Raw disk: Hard disk access at a raw, binary level, beneath the file system level.
 - Disk Appliance: A system of servers, storage nodes, and software.

- Physical tape versus Virtual tape library:
 - VTL presents a storage component (usually hard disk storage) as tape libraries or tape drives for use as storage medium with the NetWorker software.
 - Physical tape is a type of removable storage media, generally referred to as a volume or cartridge, that contains magnetic tape as its medium.
- ◆ Connectivity:
 - Local, SAN-attached:

A computer network, separate from a LAN or WAN, designed to attach shared storage devices such as disk arrays and tape libraries to servers.
 - IP-attached:

The storage device has its own unique IP address.

The component 70 percent rule

Manufacturer throughput and performance specifications based on theoretical environments are rarely, or never achieved in real backup environments. It is a best practice to never exceed 70 percent of the rated capacity of any component.

Components include:

- ◆ CPU
- ◆ Disk
- ◆ Network
- ◆ Internal bus
- ◆ Memory
- ◆ Fibre Channel

Performance and response time significantly decreases when the 70 percent utilization threshold is exceeded.

The physical tape drives, and solid state disks are the only exception to this rule, and should be used as close to 100 percent as possible. Neither the tape drives, nor solid state disks do not suffer performance degradation during heavy use.

Components of a NetWorker environment

This section describes the components of a NetWorker datazone. [Figure 3 on page 22](#) illustrates the main components in a NetWorker environment. The components and technologies that make up a NetWorker environment are listed here:

- ◆ “Datazone” on page 22
- ◆ “NetWorker Management Console” on page 22
- ◆ “NetWorker server” on page 23
- ◆ “NetWorker storage node” on page 24
- ◆ “NetWorker client” on page 25
- ◆ “NetWorker databases” on page 25
- ◆ “Optional NetWorker Application Modules” on page 26
- ◆ “Virtual environments” on page 26
- ◆ “NetWorker deduplication nodes” on page 26




Figure 3 NetWorker datazone components

Datazone

A datazone is a single NetWorker server and its client computers. Additional datazones can be added as backup requirements increase.

Note: It is recommended to have no more than 1500 clients or 3000 client instances per NetWorker datazone. This number reflects an average NetWorker server and is not a hard limit.

NetWorker Management Console

The NetWorker Management Console (NMC) is used to administer the backup server and it provides backup reporting capabilities.

The NMC often runs on the backup server, and adds significant load to the backup server. For larger environments, it is recommended to install NMC on a separate computer. A single NMC server can be used to administer multiple backup servers.

Components that determine NMC performance

Components that determine the performance of NMC are:

- ◆ TCP network connectivity to backup server: All communication between NMC and NW server is over TCP and such high-speed low-latency network connectivity is essential.
- ◆ Memory: Database tasks in larger environments are memory intensive, make sure that NMC server is equipped with sufficient memory.
- ◆ CPU: If NMC server is used by multiple users, make sure that it has sufficient CPU power to ensure that each user is given enough CPU time slices.

NetWorker server

NetWorker servers provide services to back up and recover data for the NetWorker client computers in a datazone. The NetWorker server can also act as a storage node and control multiple remote storage nodes.

Index and media management operations are some of the primary processes of the NetWorker server:

- ◆ The client file index tracks the files that belong to a save set. There is one client file index for each client.
- ◆ The media database tracks:
 - The volume name
 - The location of each saveset fragment on the physical media (file number/file record)
 - The backup dates of the save sets on the volume
 - The file systems in each save set
- ◆ Unlike the client file indexes, there is only one media database per server.
- ◆ The client file indexes and media database can grow to become prohibitively large over time and will negatively impact backup performance.
- ◆ The NetWorker server schedules and queues all backup operations, tracks real-time backup and restore related activities, and all NMC communication. This information is stored for a limited amount of time in the jobsdb which for real-time operations has the most critical backup server performance impact.

Note: The data stored in this database is not required for restore operations.

Components that determine backup server performance

Components that determine NetWorker server backup performance are:

- ◆ Minimize these system resource intensive operations on the NetWorker server during heavy loads, such as a high number of concurrent backup/clone/recover streams:
 - nsrim
 - nsrck

- ◆ The disk used to host the NetWorker server (/nsr.)

Note: The typical NetWorker server workload is from many small I/O operations. This is why disks with high latency perform poorly despite having peak bandwidth. High latency rates are the most common bottleneck of a backup server in larger environments.

- ◆ Avoid additional software layers as this adds to storage latency. For example, the antivirus software should be configured with the NetWorker databases (/nsr) in its exclusion list.
- ◆ Plan the use of replication technology carefully as it significantly increases storage latency.
- ◆ Ensure that there is sufficient CPU power for large servers to complete all internal database tasks.
- ◆ Use fewer CPUs, as systems with fewer high performance CPUs outperform systems with numerous lower performance CPUs.
- ◆ Do not attach a high number of high performance tape drives or AFTD devices directly to a backup server.
- ◆ Ensure that there is sufficient memory on the server to complete all internal database tasks.
- ◆ Off-load backups to dedicated storage nodes when possible for clients that must act as a storage node by saving data directly to backup server.

Note: The system load that results from storage node processing is significant in large environments. For enterprise environments, the backup server should backup only its internal databases (index and bootstrap).

NetWorker storage node

A NetWorker storage node can be used to improve performance by off loading from the NetWorker server much of the data movement involved in a backup or recovery operation. NetWorker storage nodes require high I/O bandwidth to manage the transfer of data transfer from local clients, or network clients to target devices.

Components that determine storage node performance

Components that determine storage node performance are:

- ◆ Performance of the target device used to store the backup.
- ◆ Connectivity of the system. For example, a storage node used for TCP network backups can save data only as fast as it is able to receive the data from clients.
- ◆ I/O bandwidth: Ensure that there is sufficient I/O bandwidth as each storage node uses available system bandwidth. Therefore, the backup performance of all devices is limited by the I/O bandwidth of the system itself.
- ◆ CPU: Ensure that there is sufficient CPU to send and receive large amounts of data.
- ◆ Do not overlap staging and backup operations with a VTL or AFTD solution by using ATA or SATA drives. Despite the performance of the array, ATA technology has significant performance degradation on parallel read and write streams.

NetWorker client

A NetWorker client computer is any computer whose data must be backed up. The NetWorker Console server, NetWorker servers, and NetWorker storage nodes are also NetWorker clients. NetWorker clients hold mission critical data and are resource intensive. Applications on NetWorker clients are the primary users of CPU, network, and I/O resources. Only read operations performed on the client do not require additional processing.

Client speed is determined by all active instances of a specific client backup at a point in time.

Components that determine NetWorker client performance

Components that determine NetWorker client performance are:

- ◆ Client backups are resource intensive operations and impact the performance of primary applications. When sizing systems for applications, be sure to consider backups and the related bandwidth requirements. Also, client applications use a significant amount of CPU and I/O resources slowing down backups.

If a NetWorker client does not have sufficient resources, both backup and application performance are negatively impacted.
- ◆ NetWorker clients with millions of files. As most backup applications are file based solutions, a lot of time is used to process all of the files created by the file system. This negatively impacts NetWorker client backup performance. For example:
 - A full backup of 5 million 20 KB files takes much longer than a backup of a half million 200 KB files, although both result in a 100 GB save set.
 - For the same overall amount of changed data, an incremental/differential backup of one thousand 100 MB files with 50 modified files takes much less time than one hundred thousand 1 MB files with 50 modified files.
- ◆ Encryption and compression are resource intensive operations on the NetWorker client and can significantly affect backup performance.
- ◆ Backup data must be transferred to target storage and processed on the backup server:
 - Client/storage node performance:
 - A local storage node: Uses shared memory and does not require additional overhead.
 - A remote storage node: Receive performance is limited by network components.
 - Client/backup server load:

Does not normally slow client backup performance unless the backup server is significantly undersized.

NetWorker databases

The factors that determine the size of NetWorker databases are available in [“NetWorker database bottlenecks” on page 32](#).

Optional NetWorker Application Modules

NetWorker Application Modules are used for specific online backup tasks. Additional application-side tuning might be required to increase application backup performance. The documentation for the applicable NetWorker module provides details.

Virtual environments

NetWorker clients can be created for virtual machines for either traditional backup or VMware Consolidated Backup (VCB). Additionally, the NetWorker software can automatically discover virtual environments and changes to those environments on either a scheduled or on-demand basis and provides a graphical view of those environments.

NetWorker deduplication nodes

A NetWorker deduplication node is an EMC Avamar[®] server that stores deduplicated backup data. The initial backup to a deduplication node should be a full backup. During subsequent backups, the Avamar infrastructure identifies redundant data segments at the source and backs up only unique segments, not entire files that contain changes. This reduces the time required to perform backups, as well as both the network bandwidth and storage space used for backups of the NetWorker Management Console.

Recovery performance factors

Recovery performance can be impeded by network traffic, bottlenecks, large files, and more. Some considerations for recovery performance are:

- ◆ File-based recovery performance depends on the performance of the backup server, specifically the client file index. Information on the client file index is available in [“NetWorker server” on page 23](#).
- ◆ The fastest method to recover data efficiently is to run multiple recover commands simultaneously by using save set recover. For example, 3 save set recover operations provide the maximum possible parallelism given the number of processes, the volume, and the save set layout.
- ◆ If multiple, simultaneous recover operations run from the same tape, be sure that the tape does not mount and start until all recover requests are ready. If the tape is used before all requests are ready, the tape is read multiple times slowing recovery performance.
- ◆ Multiplexing backups to tape slows recovery performance.


Connectivity and bottlenecks

The backup environment consists of various devices from system, storage, network, and target device components, with hundreds of models from various vendors available for each of them.

The factors affecting performance with respect to connectivity are listed here:

- ◆ Components can perform well as standalone devices, but how well they perform with the other devices on the chain is what makes the configuration optimal.
- ◆ Components on the chain are of no use if they cannot communicate to each other.
- ◆ Backups are data intensive operations and can generate large amounts of data. Data must be transferred at optimal speeds to meet business needs.
- ◆ The slowest component in the chain is considered a bottleneck.


In [Figure 4 on page 28](#), the network is unable to gather and send as much data as that of the components. Therefore, the network is the bottleneck, slowing down the entire backup process. Any single network device on the chain, such as a hub, switch, or a NIC, can be the bottleneck and slow down the entire operation.



GEN-001194

Figure 4 Network device bottleneck

As illustrated in [Figure 5 on page 29](#), the network is upgraded from a 100 base T network to a GigE network, and the bottleneck has moved to another device. The host is now unable to generate data fast enough to use the available network bandwidth. System bottlenecks can be due to lack of CPU, memory, or other resources.



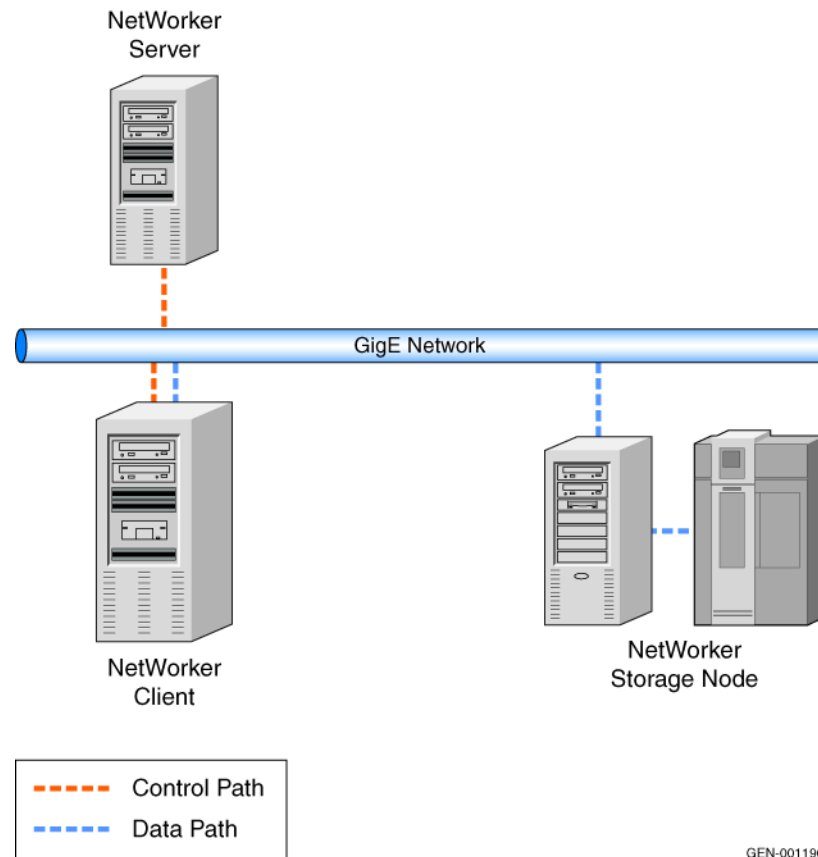
GEN-001195

Figure 5 Updated network

As illustrated in [Figure 6 on page 30](#), the NetWorker client is upgraded to a larger system to remove it as the bottleneck. With a better system and more network bandwidth, the bottleneck is now the target device. Tape devices often do not perform well as other components. Some factors that limit tape device performance are:

- ◆ Limited SCSI bandwidth
- ◆ Maximum tape drive performance reached.

Improve the target device performance by introducing higher performance tape devices, such as Fibre Channel based drives. Also, SAN environments can greatly improve performance.



GEN-001196

Figure 6 Updated client

As illustrated in [Figure 7 on page 31](#), higher performance tape devices on a SAN remove them as the bottleneck. The bottleneck device is now the storage devices. Although the local volumes are performing at optimal speeds, they are unable to use the available system, network, and target device resources. To improve the storage performance, move the data volumes to high performance external RAID arrays.




Figure 7 Dedicated SAN

Although the local volumes are performing at optimal speeds, they are unable to use the available system, network, and target device resources. To improve the storage performance, move the data volumes to high performance external RAID arrays.

As illustrated in [Figure 8 on page 32](#), the external RAID arrays have improved the system performance. The RAID arrays perform nearly as well as the other components in the chain ensuring that performance expectations are met. There will always be a bottleneck, however the impact of the bottleneck device is limited as all devices are performing at almost the same level as the other devices in the chain.




Figure 8 **Raid array**

Note: This section does not suggest that all components must be upgraded to improve performance, but attempts to explain the concept of bottlenecks, and stresses the importance of having devices that perform at similar speeds as other devices in the chain.

NetWorker database bottlenecks

This section lists factors that determine the size of NetWorker databases:

- ◆ NetWorker resource database `/nsr/res` or `networker install dir/res`: The number of configured resources.
- ◆ NetWorker jobs database (`nsr/res/jobsdb`): The number of jobs such as backups, restores, clones multiplied by number of days set for retention. This can exceed 100,000 records in the largest environments and is one of the primary performance bottlenecks.

Note: The overall size is never significant.

- ◆ For the NetWorker media database (`nsr/mm`): The number of savesets in retention and the number of labeled volumes. In the largest environments this can reach several Gigabytes of data.

- ◆ For the NetWorker client file index database (nsr/index): The number of files indexed and in the browse policy. This is normally the largest of the NetWorker databases. For storage sizing, use this formula:

$$\text{Index catalog size} = (n+(i*d))*c*160*1.5$$

where:

n = number of files to backup

d = days in cycle (time between full backups)

i = incremental data change per day in percentages

c = number of cycles online (browse policy)

The statistical average is 160 bytes per entry in the catalog.

Multiply by 1.5 to accommodate growth and error

Note: The index database can be split over multiple locations, and the location is determined on a per client bases.

Figure 9 on page 33 illustrates the overall performance degradation when the disk performance on which NetWorker media database resides is a bottleneck.

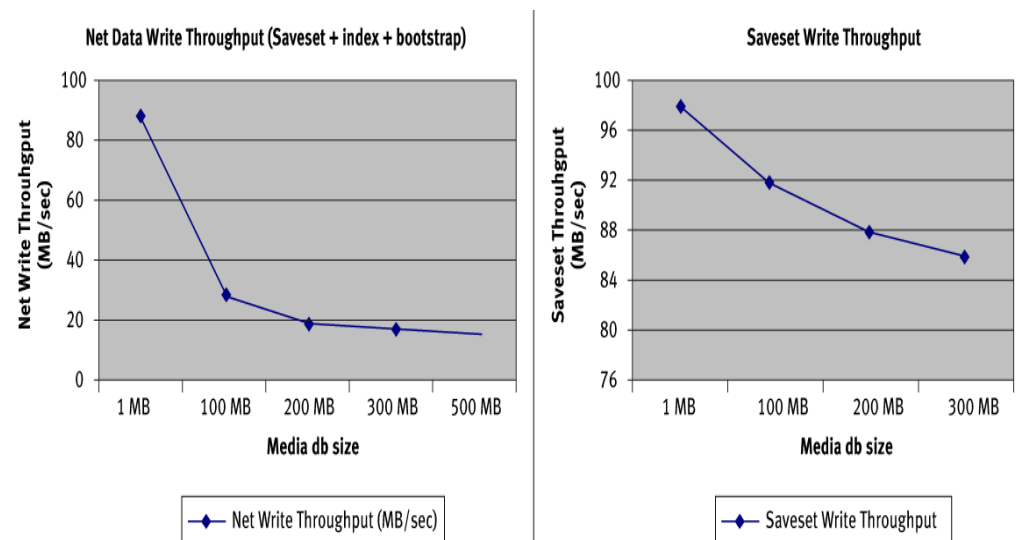


Figure 9 NetWorker server write throughput degradation

The NetWorker software has various optimization features that can be used to tune the backup environment and to optimize backup and restore performance.

This chapter includes the following topics:

- ◆ Optimize NetWorker parallelism..... 36
- ◆ Device performance tuning methods..... 38
- ◆ Network devices..... 39
- ◆ Network optimization..... 42
- ◆ Storage optimization..... 50

Optimize NetWorker parallelism

This section describes general best practices for server, group, and client parallelism.

Server parallelism

The server parallelism attribute controls how many save streams the server accepts simultaneously. The more save streams the server can accept, the faster the devices and client disks run. Client disks can run at their performance limit or the limits of the connections between them.

Server parallelism is not used to control the startup of backup jobs, but as a final limit of sessions accepted by a backup server. The server parallelism value should be as high as possible while not overloading the backup server itself.

Client parallelism

The best approach for client parallelism values is:

- ◆ For regular clients, use the lowest possible parallelism settings to best balance between the number of save sets and throughput.
- ◆ For the backup server, set highest possible client parallelism to ensure that index backups are not delayed. This ensures that groups complete as they should.

Often backup delays occur when client parallelism is set too low for the NetWorker server. The best approach to optimize NetWorker client performance is to eliminate client parallelism, reduce it to 1, and increase the parallelism based on client hardware and data configuration.

It is critical that the NetWorker server has sufficient parallelism to ensure index backups do not impede group completion.

The client parallelism values for the client that represents the NetWorker server are:

- ◆ Never set parallelism to 1
- ◆ For small environments (under 30 servers), set parallelism to at least 8
- ◆ For medium environments (31–100 servers), set parallelism to at least 12
- ◆ For larger environments (100+ servers), set parallelism to at least 16

These recommendations assume that the backup server is a dedicated backup server. The backup server should always be a dedicated server for optimum performance.

Group parallelism

The best approach for group parallelism values is:

- ◆ Create save groups with a maximum of 50 clients with group parallelism enforced. Large save groups with more than 50 clients can result in many operating system processes starting at the same time causing temporary operating system resource exhaustion.
- ◆ Stagger save group start times by a small amount to reduce the load on the operating system. For example, it is best to have 4 save groups, each with 50 clients, starting at 5 minute intervals than to have 1 save group with 200 clients.

Multiplexing

The Target Sessions attribute sets the target number of simultaneous save streams that write to a device. This value is not a limit, therefore a device might receive more sessions than the Target Sessions attribute specifies. The more sessions specified for Target Sessions, the more save sets that can be multiplexed (or interleaved) onto the same volume.

[“AFTD device target and max sessions” on page 40](#) provides additional information on device Target Sessions.

Performance tests and evaluation can determine whether multiplexing is appropriate for the system. Follow these guidelines when evaluating the use of multiplexing:

- ◆ Find the maximum rate of each device. Use the bigasm test described in [“The bigasm directive” on page 56](#).
- ◆ Find the backup rate of each disk on the client. Use the uasm test described in [“The uasm directive” on page 56](#).

If the sum of the backup rates from all disks in a backup is greater than the maximum rate of the device, do not increase server parallelism. If more save groups are multiplexed in this case, backup performance will not improve, and recovery performance might slow down.

Device performance tuning methods

These sections address specific device-related areas that can improve performance.

Input/output transfer rate

Input/output (I/O) transfer rates can affect device performance. The I/O rate is the rate at which data is written to a device. Depending on the device and media technology, device transfer rates can range from 500 KB per second to 200 MB per second. The default block size and buffer size of a device affect its transfer rate. If I/O limitations interfere with the performance of the NetWorker server, try upgrading the device to affect a better transfer rate.

Built-in compression

Turn on device compression to increase effective throughput to the device. Some devices have a built-in hardware compression feature. Depending on how compressible the backup data is, this can improve effective data throughput, from a ratio of 1.5:1 to 3:1.

Drive streaming

To obtain peak performance from most devices, stream the drive at its maximum sustained throughput. Without drive streaming, the drive must stop to wait for its buffer to refill or to reposition the media before it can resume writing. This can cause a delay in the cycle time of a drive, depending on the device.

Device load balancing

Balance data load for simultaneous sessions more evenly across available devices by adjusting target and max sessions per device. This parameter specifies the minimum number of save sessions to be established before the NetWorker server attempts to assign save sessions to another device. More information on device target and max sessions is available at [“AFTD device target and max sessions” on page 40](#).

Network devices

If data is backed up from remote clients, the routers, network cables, and network interface cards affect the backup and recovery operations. This section lists the performance variables in network hardware, and suggests some basic tuning for networks. The following items address specific network issues:

- ◆ Network I/O bandwidth:

The maximum data transfer rate across a network rarely approaches the specification of the manufacturer because of network protocol overhead.

Note: The following statement concerning overall system sizing must be considered when addressing network bandwidth.

Each attached tape drive (physical VTL or AFTD) uses available I/O bandwidth, and also consumes CPU as data still requires processing.

- ◆ Network path:

Networking components such as routers, bridges, and hubs consume some overhead bandwidth, which degrades network throughput performance.

- ◆ Network load:

- Do not attach a large number of high-speed NICs directly to the NetWorker server, as each IP address use significant amounts of CPU resources. For example, a mid-size system with four 1 GB NICs uses more than 50 percent of its resources to process TCP data during a backup.
- Other network traffic limits the bandwidth available to the NetWorker server and degrades backup performance. As the network load reaches a saturation threshold, data packet collisions degrade performance even more.

DataDomain

Backup to DataDomain storage can be configured by using multiple technologies:

- ◆ Backup to VTL:

NetWorker devices are configured as tape devices and data transfer occurs over Fiber Channel.

Information on VTL optimization is available in [“Number of virtual device drives versus physical device drives” on page 41](#).

- ◆ Backup to AFTD over CIFS or NFS:

- Overall network throughput depends on the CIFS and NFS performance which depends on network configuration.

[“Network optimization” on page 42](#) provides best practices on backup to AFTD over CIFS or NFS.

- Inefficiencies in the underlying transport limits backup performance to 70-80 percent of the link speed. For optimal performance, NetWorker release 7.5 Service Pack 2 or later is required.

- ◆ Backup to DataDomain by using native device type:

- NetWorker 7.6 Service Pack 1 provides a new device type designed specifically for native communication to DataDomain storage over TCP/IP links.

- With proper network optimization, this protocol is capable of using up to 95 percent of the link speed even at 10 Gb/sec rates and is currently the most efficient network transport.
- Each DataDomain device configured in NetWorker is limited to a maximum of 10 parallel backup streams. If higher parallelism is required, configure more devices to a limit defined by the NetWorker server edition.

Note: Despite the method used for backup to DataDomain storage, the aggregate backup performance is limited by the maximum ingress rate of the specific DataDomain model.

- ◆ The minimum required memory for a NetWorker DataDomain-OST device with each device total streams set to 10 is approximately 160 MB. Each OST stream for BOOST takes an additional 16 MB of memory .

AFTD device target and max sessions

This section describes for all supported operating systems, the optimal Advanced File Type Device (AFTD) device target, and max sessions settings for the NetWorker software. Details for NetWorker versions 7.6 and earlier, and 7.6 Service Pack 1 and later software are included.

NetWorker 7.6 and earlier software

The current NetWorker 7.6 and earlier default settings for AFTD target sessions (4) and max sessions (512) are not optimal for AFTD performance.

To optimize AFTD performance for NetWorker 7.6 and earlier, change the default values:

- ◆ Set device target sessions from 4 to 1.
- ◆ Set device max sessions from 512 to 32 to avoid disk thrashing.

NetWorker 7.6 Service Pack 1 and later

The defaults for AFTD target sessions and max device sessions are now set to the optimal values for AFTD performance:

- ◆ Device target sessions is 1
- ◆ Device max sessions is 32 to avoid disk thrashing

If required, both Device target, and max session attributes can be modified to reflect values appropriate for the environment.

Number of virtual device drives versus physical device drives

The following is based on the 70 percent utilization of a Fibre Channel port:

- ◆ For LTO-3: 3 virtual devices for every 2 physical devices planned.
- ◆ For LTO-4: 3 virtual devices for each physical device planned.

On the port:

The performance of each of these tape drives on the same port degrades with the number of attached devices. For example:

- ◆ If the first virtual drive reaches the 150 MB per second limit.
- ◆ The second virtual drive will not exceed 100 MB per second.
- ◆ The third virtual drive will not exceed 70 MB per second.

Network optimization

This section explains the following:

- ◆ [“Advanced configuration optimization” on page 42](#)
- ◆ [“Operating system TCP stack optimization” on page 42](#)
- ◆ [“Advanced tuning” on page 43](#)
- ◆ [“Expected NIC throughput values” on page 43](#)
- ◆ [“Network latency” on page 43](#)
- ◆ [“Ethernet duplexing” on page 45](#)
- ◆ [“Firewalls” on page 45](#)
- ◆ [“Jumbo frames” on page 45](#)
- ◆ [“Congestion notification” on page 45](#)
- ◆ [“TCP buffers” on page 46](#)
- ◆ [“NetWorker socket buffer size” on page 47](#)
- ◆ [“IRQ balancing and CPU affinity” on page 47](#)
- ◆ [“Interrupt moderation” on page 48](#)
- ◆ [“TCP offloading” on page 48](#)
- ◆ [“Name resolution” on page 49](#)

Advanced configuration optimization

The EMC Technical Note, *Configuring TCP Networks and Network Firewalls for EMC NetWorker* provides instructions on advanced configuration options such as multihomed systems, trunking, and so on.

The default TCP operating system parameters are tuned for maximum compatibility with legacy network infrastructures, but not for maximum performance.

Operating system TCP stack optimization

The common rules for optimizing the operating system TCP stack for all use cases are listed here:

- ◆ Disable software flow control.
- ◆ Increase TCP buffer sizes.
- ◆ Increase TCP queue depth.
- ◆ Use PCIeExpress for 10 GB NICs. Other I/O architectures do not have enough bandwidth.

More information on PCIeExpress is available in [“PCI-X and PCIeExpress considerations:” on page 18](#).

Rules that depend on environmental capabilities are listed here:

- ◆ Some operating systems have internal auto-tuning of the TCP stack. This produces good results in a non-heterogeneous environment. However, for heterogeneous, or routed environments disable TCP auto-tuning.
- ◆ Enable jumbo frames when possible.

Note: It is required that *all* network components in the data path are able to handle jumbo frames. Do not enable jumbo frames if this is not the case.

- ◆ TCP hardware offloading is beneficial if it works properly. However it can cause CRC mis-matches. Be sure to monitor for errors if it is enabled.
- ◆ TCP windows scaling is beneficial if it is supported by all network equipment in the chain.
- ◆ TCP congestion notification can cause problems in heterogeneous environments. Only enable it in single operating system environments.

Advanced tuning

IRQ processing for high-speed NICs is very expensive, but can provide enhanced performance by selecting specific CPU cores. Specific recommendations depend on the CPU architecture.

Expected NIC throughput values

Common NIC throughput values are in the following ranges:

- ◆ 100 Mb link = 6–8 MB/s
- ◆ 1 Gb link = 45–65MB/s
- ◆ 10 Gb link = 150–350 Mb/s

With optimized values, throughput for high-speed links can be increased to the following:

- ◆ 100 Mb link = 12 MB/s
- ◆ 1 Gb link = 110MB/s
- ◆ 10 Gb link = 1100 MB/s

The Theoretical maximum throughput for a 10 Gb Ethernet link is 1.164 GB/s per direction calculated by converting bits to bytes and removing the minimum Ethernet, IP and TCP overheads.

Network latency

Increased network TCP latency has a negative impact on overall throughput, despite the amount of available link bandwidth. Longer distances or more hops between network hosts can result in lower overall throughput.

Network latency has a high impact on the efficiency of bandwidth use.

For example, [Figure 10 on page 44](#) and [Figure 11 on page 44](#) illustrate backup throughput on the same network link, with varying latency.

Note: For these examples, non-optimized TCP settings were used.

Backup Throughput 10/100 MBps




Figure 10 Network latency on 10/100 MB per second

Backup Throughput on 1GIG

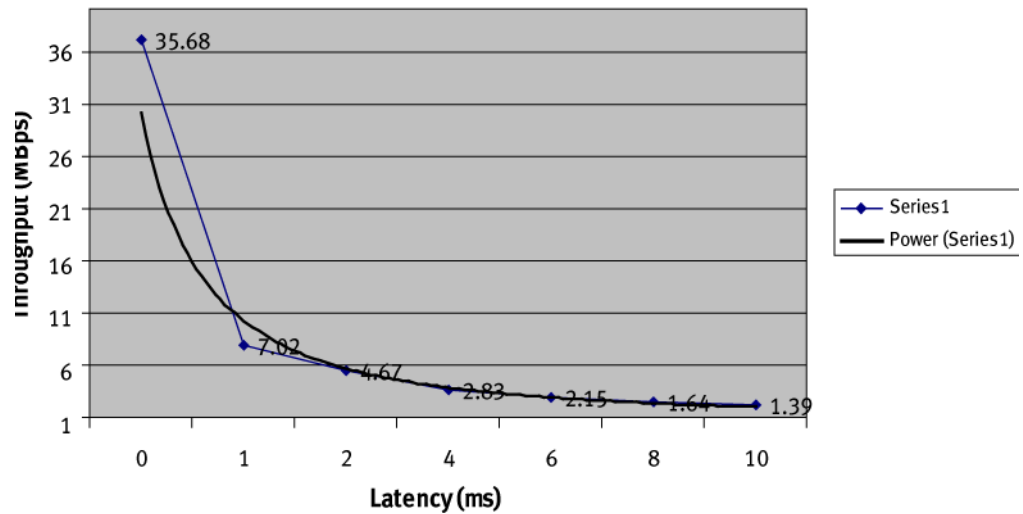


Figure 11 Network latency on 1 GIG

Ethernet duplexing

Network links that perform in half-duplex mode cause decreased NetWorker traffic flow performance. For example, a 100 Mb half-duplex link results in backup performance of less than 1 MB/s.

The default configuration setting on most operating systems for duplexing is auto negotiated as recommended by IEEE802.3. However, auto negotiation requires that the following conditions are met:

- ◆ Proper cabling
- ◆ Compatible NIC adapter
- ◆ Compatible switch

Auto negotiation can result in a link performing as half-duplex.

To avoid issues with auto negotiation, force full-duplex settings on the NIC. Forced full-duplex setting must be applied to both sides of the link. Forced full-duplex on only one side of the link results in failed auto negotiation on the other side of the link.

Firewalls

The additional layer on the I/O path in a hardware firewall increases network latency, and reduces the overall bandwidth use.

It is recommended to avoid using software firewalls on the backup server as it processes a large number of packets resulting in significant overhead.

Details on firewall configuration and impact are available in the technical note, *Configuring TCP Networks and Network Firewalls for EMC NetWorker*.

Jumbo frames

It is recommended to use jumbo frames in environments capable of handling them.

If both the source and the computers, and all equipment in the data path are capable of handling jumbo frames, increase the MTU to 9 KB:

These examples are for Linux and Solaris operating systems:

- ◆ Linux: `ifconfig eth0 mtu 9000 up`
- ◆ Solaris: `nxge0 accept-jumbo 1`

Congestion notification

This section describes how to disable congestion notification algorithms.

- ◆ Windows 2008 R2 only:
 1. Disable optional congestion notification algorithms:
`C:\> netsh interface tcp set global ecncapability=disabled`
 2. Advanced TCP algorithm provides the best results on Windows. However disable advanced TCP algorithm if both sides of the network conversion are not capable of the negotiation:
`C:\> netsh interface tcp set global congestionprovider=ctcp`

- ◆ Linux:
 1. Check for non-standard algorithms:
`cat /proc/sys/net/ipv4/tcp_available_congestion_control`
 2. Disable ECN:
`echo 0 >/proc/sys/net/ipv4/tcp_ecn`
- ◆ Solaris:
Disable TCP Fusion if present:
`set ip:do_tcp_fusion = 0x0`

TCP buffers

For high-speed network interfaces, increase size of TCP send/receive buffers:

- ◆ Linux:


```
echo 262144 >/proc/sys/net/core/rmem_max
echo 262144 >/proc/sys/net/core/wmem_max
echo 262144 >/proc/sys/net/core/rmem_default
echo 262144 >/proc/sys/net/core/wmem_default
echo '8192 524288 2097152' >/proc/sys/net/ipv4/tcp_rmem
echo '8192 524288 2097152' >/proc/sys/net/ipv4/tcp_wmem
```

Set the recommended RPC value:

```
sunrpc.tcp_slot_table_entries = 64
```

Another method is to enable dynamic TCP window scaling. This requires compatible equipment in the data path:

```
sysctl -w net.ipv4.tcp_window_scaling=1
```
- ◆ Solaris:


```
tcp_max_buf 10485760
tcp_cwnd_max 10485760
tcp_recv_hiwat 65536
tcp_xmit_hiwat 65536
```
- ◆ AIX

Modify the values for the parameters in /etc/rc.net if the values are lower than the recommended. The number of bytes a system can buffer in the kernel on the receiving sockets queue:

```
no -o tcp_recvspace=524288
```

The number of bytes an application can buffer in the kernel before the application is blocked on a send call:

```
no -o tcp_sendspace=524288
```

- ◆ Windows:
 - The default buffer sizes maintained by the Windows operating system are sufficient.
 - Set the registry entry:
 - AdditionalCriticalWorkerThreads: DWORD=10**
 - If the NIC drivers are able to create multiple buffers or queues at the driver-level, enable it at the driver level. For example, Intel 10 Gb NIC drivers by default have RSS Queues set to 2, and the recommended value for optimum performance is 16.
 - The Windows 2008 server introduces a method to auto tune the TCP stack. If a server on the LAN or a network device in the datazone such as a router or switch does not support TCP Windows scaling, backups can fail. To avoid failed backups, and ensure optimal NetWorker operations, apply the Microsoft Hotfix KB958015 to the Windows 2008 Server, and set the auto tuning level value to highlyrestricted:
 1. Check the current TCP settings:
 - C:\> netsh interface tcp show global**
 2. If required, restrict the Windows TCP receive side scaling auto tuning level:
 - C:\> netsh interface tcp set global autotuninglevel=highlyrestricted**

Note: If the hotfix KB958015 is not applied, the autotuning level must be set to disabled rather than highlyrestricted.

NetWorker socket buffer size

To force the use of a larger TCP send/receive window from NetWorker, include these in the NetWorker start script:

```
NSR SOCK BUF SIZE=65536
export NSR SOCK BUF SIZE
```

- ◆ The optimal TCP socket buffer for a 1 GB network is 64 KB.
- ◆ The optimal TCP socket buffer for a 10 GB network is 256KB. Include this in the NetWorker start script:

```
NSR SOCK BUF SIZE=262144
```

IRQ balancing and CPU affinity

A high-speed network interface that uses either multiple 1 Gb interfaces or one 10 Gb interface benefits from disabled IRQ balancing and binding to specific CPU core processing.

Note: The general rule is that only one core per physical CPU should handle NIC interrupts. Use multiple cores per CPU only if there are more NICs than CPUs. However, transmitting and receiving should always be handled by the same CPU without exception.

These examples are for Linux and Solaris operating systems:

- ◆ Linux:

1. Disable IRQ balancing and set CPU affinity manually:

```
service irqbalance stop
chkconfig irqbalance off
```

2. Tune the CPU affinity for the eth0 interface:

```
grep eth0 /proc/interrupts
```

3. Tune the affinity for the highest to the lowest. For example:

```
echo 80 > /proc/irq/177/smp_affinity
echo 40 > /proc/irq/166/smp_affinity
```

Note: SMP affinity works only for IO-APIC enabled device drivers. Check for the IO-APIC capability of a device by using `cat /proc/interrupts`, or by referencing the device documentation.

- ◆ Solaris:

Interrupt only one core per CPU. For example, for a system with 4 CPUs and 4 cores per CPU, use this command:

```
psradm -i 1-3 5-7 9-11 13-15
```

Additional tuning depends on the system architecture.

These are examples of successful settings on a Solaris system with a T1/T2 CPU (Niagara):

```
ddi_msix_alloc_limit 8
tcp_queue_wput 1
ip_soft_rings_cnt 64
ip_queue_fanout 1
```

Some NIC drivers artificially limit interrupt rates to reduce peak CPU use. However, this also limits the maximum achievable throughput. If a NIC driver is set for "Interrupt moderation," disable it for optimal network throughput.

Interrupt moderation

On Windows, for a 10GB network, it is recommended to disable interrupt moderation for the network adapter to improve network performance.

TCP offloading

For systems with NICs capable of handling TCP packets at a lower level, enable TCP offloading on the operating system to:

- ◆ Increase overall bandwidth utilization
- ◆ Decrease the CPU load on the system

Note: Not all NICs that market offloading capabilities are fully compliant with the standard.

- ◆ For a Windows 2008 server, use this command to enable TCP offloading:

```
C:\> netsh interface tcp set global chimney=enabled
```

- ◆ For a Windows 2008 R2 server, use these commands with additional properties to enable TCP offloading:

```
C:\> netsh interface tcp set global dca=enabled  
C:\> netsh interface tcp set global netdma=enabled
```

- ◆ Disable TCP offloading for older generation NIC cards that exhibit problems such as backup sessions that hang, or fail with RPC errors similar to this:

```
Connection reset by peer
```

Name resolution

The NetWorker server relies heavily on the name resolution capabilities of the operating system.

For a DNS server, set low-latency access to the DNS server to avoid performance issues by configuring, either of these:

- ◆ Local DNS cache

or

- ◆ Local non-authoritative DNS server with zone transfers from the main DNS server

Ensure that the server name and hostnames assigned to each IP address on the system are defined in the hosts file to avoid DNS lookups for local hostname checks.

Storage optimization

This section describes settings for NetWorker server and storage node disk optimization.

NetWorker server and storage node disk write latency

This section describes requirements for NetWorker server and storage node write latency.

Write latency for /nsr on NetWorker servers, and storage nodes is more critical for the storage hosting /nsr than is the overall bandwidth. This is because NetWorker uses very large number of small random I/O for internal database access. [Table 3 on page 50](#) lists the effects on performance for disk write latency during NetWorker backup operations.

Table 3 Disk write latency results and recommendations

Disk write latency in milliseconds (ms)	Effect on performance	Recommended
25 ms and below	<ul style="list-style-type: none"> Stable backup performance Optimal backup speeds 	Yes
50 ms	<ul style="list-style-type: none"> Slow backup performance (the NetWorker server is forced to throttle database updates) Delayed & failed NMC updates 	No
100 ms	Failed savegroups and sessions	No
150–200 ms	<ul style="list-style-type: none"> Delayed NetWorker daemon launch Unstable backup performance Unprepared volumes for write operations Unstable process communication 	No

Note: Avoid using synchronous replication technologies or any other technology that adversely impacts latency.

Recommended server and storage node disk settings

This section lists recommendations for optimizing NetWorker server and storage node disk performance:

- ◆ For NetWorker servers under increased load (number of parallel sessions occurring during a backup exceeds 100 sessions), dedicate a fast disk device to host NetWorker databases.
- ◆ For disk storage configured for the NetWorker server, use RAID-10.
- ◆ For large NetWorker servers with server parallelism higher than 400 parallel sessions, split the file systems used by the NetWorker server. For example, split the /nsr folder from a single mount to multiple mount points for:

```

/nsr
/nsr/res
/nsr/index
/nsr/mm

```

- ◆ For NDMP backups on the NetWorker server, use a separate location for /nsr/tmp folder to accommodate large temporary file processing.
- ◆ Use the operating system to handle parallel file system I/O even if all mount points are on the same physical location. The operating system handles parallel file system I/O more efficiently than the NetWorker software.
- ◆ Use RAID-3 for disk storage for AFTD.
- ◆ For antivirus software, disable scanning of the NetWorker databases. If the antivirus software is able to scan the /nsr folder, performance degradation, time-outs, or NetWorker database corruption can occur because of frequent file open/close requests. The antivirus exclude list should also include NetWorker storage node locations used for Advanced File Type Device (AFTD).

Note: Disabled antivirus scanning of specific locations might not be effective if it includes all locations during file access, despite the exclude list if it skips scanning previously accessed files. Contact the specific vendor to obtain an updated version of the antivirus software.

- ◆ For file caching, aggressive file system caching can cause commit issues for:
 - The NetWorker server: all NetWorker databases can be impacted (nsr\res, nsr\index, nsr\mm).
 - The NetWorker storage node: When configured to use Advanced File Type Device (AFTD).
Be sure to disable delayed write operations, and use file system driver Flush and Write-Through commands instead.
- ◆ Disk latency considerations for the NetWorker server are higher than for typical server applications as NetWorker utilizes committed I/O: Each write to NW internal database must be acknowledged and flushed before next write is attempted. This is to avoid any potential data loss in internal databases. These are considerations for /nsr in cases where storage is replicated or mirrored:
 - Do not use software based replication as it adds an additional layer to I/O throughput and causes unexpected NetWorker behavior.
 - With hardware based replication, the preferred method is asynchronous replication as it does not add latency on write operations.
 - Do not use synchronous replication over long distance links, or links with non-guaranteed latency.
 - SANs limit local replication to 12 km and longer distances require special handling.
 - Do not use TCP networks for synchronous replication as they do not guarantee latency.
 - Consider the number of hops as each hardware component adds latency.

This chapter describes how to test and understand bottlenecks by using available tools including NetWorker programs such as bigasm and uasm. This chapter includes the following topics:

- ◆ Determine symptoms 52
- ◆ Monitor performance..... 53
- ◆ Determine bottlenecks by using a generic FTP test..... 54
- ◆ Test the performance of the setup by using dd..... 55
- ◆ Test disk performance by using bigasm and uasm 56

Determine symptoms

Considerations for determining the reason for poor backup performance are listed here:

- ◆ Is the performance consistent for the entire duration of the backup?
- ◆ Do the backups perform better when started at a different time?
- ◆ Is it consistent across all save sets for the clients?
- ◆ Is it consistent across all clients with similar system configuration using a specific storage node?
- ◆ Is it consistent across all clients with similar system configuration in the same subnet?
- ◆ Is it consistent across all clients with similar system configuration and applications?

Observe how the client performs with different parameters. Inconsistent backup speed can indicate problems with software or firmware.

For each NetWorker client, answer these questions:

- ◆ Is the performance consistent for the entire duration of the backup?
- ◆ Is there a change in performance if the backup is started at a different time?
- ◆ Is it consistent across all clients using specific storage node?
- ◆ Is it consistent across all save sets for the client?
- ◆ Is it consistent across all clients in the same subnet?
- ◆ Is it consistent across all clients with similar operating systems, service packs, applications?
- ◆ Does the backup performance improve during the save or does it decrease?

These and similar questions can help to identify the specific performance issues.

Monitor performance

Monitor the I/O, disk, CPU, and network performance by using native performance monitoring tools such as:

- ◆ Windows: Perfmon
- ◆ UNIX: iostat, vmstat, or netstat commands

Unusual activity before, during, and after backups can determine that devices are using excessive resources.

By using these tools to observe performance over a period of time, resources consumed by each application, including NetWorker are clearly identified.

If it is discovered that slow backups are due to excessive network use by other applications, this can be corrected by changing backup schedules.

Note: High CPU use is often the result of waiting for external I/O, not insufficient CPU power. This is indicated by high CPU use inside SYSTEM versus user space.

On Windows, if a lot of time is spent on Deferred Procedure Calls, it often indicates a problem with device drivers.

Determine bottlenecks by using a generic FTP test

Without using NetWorker components, determine whether the bottleneck is in the network or the tape device by using a generic FTP test:

1. Create a large data file on the NetWorker client and send it to the storage node by using FTP.
2. Make note of the time it takes for the file to transfer.
3. Compare the time noted in [step 2](#) with current backup performance:
 - If the ftp performs much faster than the backups, then the bottleneck might be with the tape devices.
 - If the ftp performs at a similar rate, then the bottleneck might be in the network.
4. Compare results by using active FTP versus passive FTP transfer. NetWorker backup performance is greatly impacted by the capabilities of the underlying network and the network packets used by the NetWorker software.

If there is large difference in the transfer rate, or one type of FTP transfer has spikes, it might indicate the presence of network components that perform TCP packet re-assembly. This causes the link to perform in half-duplex mode, despite all physical parts that are in full-duplex mode.

Note: Do *not* use local volumes to create and transfer files for ftp tests, use backup volumes.

Test the performance of the setup by using dd

Without using NetWorker components, use the generic dd test to compare device throughput to the manufacturer's suggested throughput:

1. Create a large data file on the storage node and use dd to send it to the target device:

```
date; dd if=/tmp/5GBfile of=/dev/rmt/0cbn bs= 1MB; date
```

2. Make note of the time it takes for the file to transfer, and compare it with the current tape performance.

Test disk performance by using bigasm and uasm

The bigasm and uasm directives are NetWorker based tests used to verify performance.

The bigasm directive

The bigasm directive generates a specific sized file, and transfers the file over a network or a SCSI connection. The file is then written to a tape or another target device. The bigasm directive creates a stream of bytes in memory and saves them to the target device that eliminates disk access. This helps to test the speed of NetWorker clients, network, and the tape devices ignoring disk access.

Create a bigasm directive to generate a very large save set.

The bigasm directive ignores disk access to test the performance of client, network and tape.

The uasm directive

The uasm directive reads from the disk at maximum speeds to identify disk based bottlenecks. For example:

```
uasm -s filename > NUL
```

The uasm directive tests disk read speeds, and by writing data to a null device can identify disk-based bottlenecks.