

# Novell NetWare® 6.5

[www.novell.com](http://www.novell.com)

February 28, 2005

SERVER MEMORY ADMINISTRATION  
GUIDE



**Novell®**

## Legal Notices

Novell<sup>®</sup>, Inc. makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc. makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc. reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

You may not use, export, or re-export this product in violation of any applicable laws or regulations including, without limitation, U.S. export regulations or the laws of the country in which you reside.

Copyright © 1993-2005 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

Novell, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.novell.com/company/legal/patents/> and one or more additional patents or pending patent applications in the U.S. and in other countries.

Novell, Inc.  
404 Wyman Street, Suite 500  
Waltham, MA 02451  
U.S.A.

[www.novell.com](http://www.novell.com)

NetWare 6.5 Server Memory Administration Guide  
[February 28, 2005](#)

**Online Documentation:** To access the online documentation for this and other Novell products, and to get updates, see [www.novell.com/documentation](http://www.novell.com/documentation).

## **Novell Trademarks**

eDirectory is a trademark of Novell, Inc.

NetWare is a registered trademark of Novell, Inc. in the United States and other countries.

NetWare Loadable Module and NLM are trademarks of Novell, Inc.

Novell is a registered trademark of Novell, Inc. in the United States and other countries.

Novell Storage Services is a trademark of Novell, Inc.

Novell Technical Services is a service mark of Novell, Inc.

## **Third-Party Materials**

All third-party products are the property of their respective owners.



# Contents

- About This Guide** **7**
  
- 1 Overview of Server Memory** **9**
  - Introduction to NetWare Memory . . . . . 9
  - Caching . . . . . 9
    - Traditional File System . . . . . 10
    - Novell Storage Services . . . . . 10
    - Directory Buffers and File Cache Buffers . . . . . 11
    - Garbage Collection . . . . . 11
  - Logical Memory Addressing . . . . . 12
  - Protected Address Spaces . . . . . 12
    - Protected Address Spaces and Logical Memory Addressing . . . . . 13
    - Modules Not Allowed in Protected Address Spaces . . . . . 13
  - Virtual Memory . . . . . 14
  - What's Next . . . . . 15
  
- 2 What's New** **17**
  - NetWare 6.5 Support Pack 3 Release . . . . . 17
  - NetWare 6.5 Support Pack 2 Release . . . . . 17
  - NetWare 6.5 Initial Release . . . . . 17
  
- 3 Managing Server Memory** **19**
  - Using Server Memory Management Tools . . . . . 19
    - Using Novell Remote Manager . . . . . 19
    - Server Memory Pool Summary Page . . . . . 36
    - Using Set Parameters . . . . . 38
    - Using Command Line Commands . . . . . 38
  - Using Protected Address Spaces . . . . . 38
  - Loading Protected Address Spaces . . . . . 39
  - Unloading Protected Address Spaces . . . . . 41
  - Viewing Information about Address Spaces . . . . . 42
  - What's Next . . . . . 42
  
- 4 Optimizing Server Memory** **43**
  - Assessing Server RAM . . . . . 43
    - Traditional File System . . . . . 43
    - Novell Storage Services . . . . . 44
  - Ensuring That Devices Can Access Memory . . . . . 44
  - Determining Which Processes Use Too Much Memory . . . . . 44
  - Tuning Server Memory . . . . . 45
    - Tuning File Cache . . . . . 45
    - Tuning Logical Address Spaces . . . . . 49
    - Tuning Directory Cache for the Traditional File System . . . . . 50
  - Using Virtual Memory . . . . . 52
    - Using Swap Files . . . . . 53
    - Responding to Disk Thrashing . . . . . 54

Optimizing Garbage Collection . . . . .	55
---	----

<b>A Documentation Updates</b>	<b>57</b>
--------------------------------	-----------

# About This Guide

This guide describes how memory works on a NetWare<sup>®</sup> 6.5 server and how to optimize and manage that server memory. This guide is intended for network administrators and is divided into the following sections:

- ♦ Chapter 1, “Overview of Server Memory,” on page 9
- ♦ Chapter 2, “What’s New,” on page 17
- ♦ Chapter 3, “Managing Server Memory,” on page 19
- ♦ Chapter 4, “Optimizing Server Memory,” on page 43

## Additional Documentation

For more documentation on managing server memory, see the *NetWare 6.5 Novell Remote Manager Administration Guide for NetWare*.

## Documentation Updates

For the most recent version of the *Server Memory Administration Guide*, see the [NetWare 6.5 Documentation Web site \(http://www.novell.com/documentation/nw65/smem\\_enu/data/front.html\)](http://www.novell.com/documentation/nw65/smem_enu/data/front.html).

## Documentation Conventions

In this documentation, a greater-than symbol (>) is used to separate actions within a step and items within a cross-reference path.

A trademark symbol (<sup>®</sup>, <sup>™</sup>, etc.) denotes a Novell<sup>®</sup> trademark. An asterisk (\*) denotes a third-party trademark.

When a single pathname can be written with a backslash for some platforms or a forward slash for other platforms, the pathname is presented with a backslash. Users of platforms that require a forward slash, such as Linux\* and UNIX\*, should use forward slashes as required by your software.

## User Comments

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comment feature at the bottom of each page of the online documentation, or go to [www.novell.com/documentation/feedback.html](http://www.novell.com/documentation/feedback.html) and enter your comments there.





# 1

## Overview of Server Memory

The NetWare® server operating system provides increased memory efficiency with three important features: logical memory addressing, virtual memory, and protected memory. This section is a general introduction to these and other aspects of the NetWare memory architecture.

To manage memory and virtual memory, you need to understand the following:

- ◆ [Introduction to NetWare Memory \(page 9\)](#)
- ◆ [Logical Memory Addressing \(page 12\)](#)
- ◆ [Protected Address Spaces \(page 12\)](#)
- ◆ [Virtual Memory \(page 14\)](#)

### Introduction to NetWare Memory

NetWare 6.5 requires a minimum of 512 MB of RAM and can access up to 64 GB of RAM. After the server.exe file is loaded, the remaining memory below 4 GB is assigned to “**cache memory**” and virtual memory, which is available for NetWare Loadable Module™ (NLM™) programs and other processes to use.

**NOTE:** The maximum memory that your server requires depends on the products and components that you install on your server and the number of applications that you run on your server that use protected memory.

The memory allocator cannot allocate memory above 4 GB. If you install more than 4 GB, this memory is used by the virtual memory system, NSS, and applications such as eDirectory™. The traditional file system does not use memory above 4 GB.

The applications that run in the memory above 4 GB leave more memory below 4 GB for the memory allocator to use. So, while you cannot configure the memory above 4 GB, adding more memory has the following performance advantages:

- ◆ More memory is available to those systems and applications that run in protected memory or in the space above 4 GB.
- ◆ The more available memory below 4 GB the memory allocator has to use, then the faster the server performs.

### Caching

A cache is a temporary, quickly accessible area in RAM that stores frequently used information from the hard disk. NetWare uses caching to increase the speed of reads and writes on server hardware. As long as a block of data stays in RAM, NetWare can access it much faster than accessing data on the disk.

A *cache hit* means that the server found the data in RAM rather than having to go to disk. This provides a large performance increase.

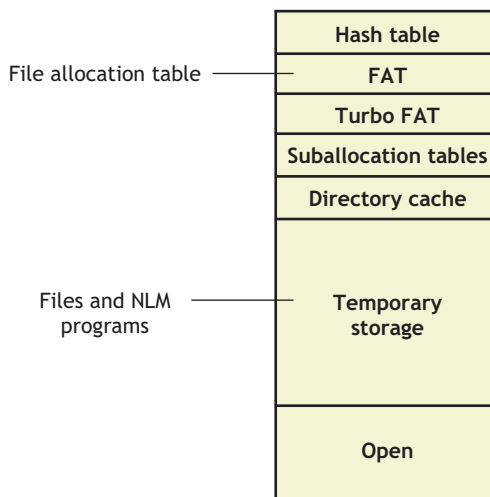
Having sufficient RAM to cache disk data results in a high percentage of cache hits. The more NLM programs and other services you load, the less RAM the server has available for caching data. This lowers your cache hit percentage and slows network performance. For information on disk cache utilization, see [“Tuning Server Memory” on page 45](#).

When data is stored in cache, it is stored in 4 KB blocks called cache buffers. Recently used data is kept in cache buffers in case it is needed again.

## Traditional File System

Cache memory is used for all the server's processing needs, such as the following:

- ◆ Storing volume File Allocation Tables (FAT) and suballocation tables
- ◆ Storing volume directory entry tables
- ◆ Storing user data files
- ◆ Storing NLM programs, such as LAN drivers, disk drivers, and management utilities
- ◆ Building hash tables for directory names
- ◆ Creating protected address spaces
- ◆ Allocated memory used by NLM programs
- ◆ Name cache and file cache



Memory from cache is allocated to an NLM program when requested. When an NLM program is finished using the memory, it releases the memory and returns it to cache. NetWare also provides a [“garbage collection”](#) process to recover unused memory.

## Novell Storage Services

In NetWare 6.5, Novell<sup>®</sup> Storage Services™ (NSS) is always loaded on the server in order to run and cache memory is always allocated. NSS uses memory differently than the traditional file system. Some of these differences are the following:

- ◆ NSS caches files larger than 128 KB in the memory above 4 GB. This improves performances for these files and for smaller files because they do not compete for the same memory.

- ◆ NSS has its own cache blocks, although they are counted as total free blocks in the system. The memory allocator can use them for applications other than NSS.
- ◆ By default, NSS maintains the Closed File Cache Size parameter at 100,000. Each cache is about 1 KB, so 100 MB of memory is eventually dedicated to Closed File Cache Size. NSS does not automatically allocate 100 MB of memory for this parameter; it waits for 100,000 files to close before it reaches 100 MB.

If the server detects that it needs more memory, some of this 100 MB will automatically be sent to the free pools. NSS adjusts itself as pressure on the memory system increases. For more information about this parameter, see [“Cache Hit Percentage” on page 46](#).

You must use specific parameters for tuning your server when using NSS. For more information, see [“Tuning File Cache” on page 45](#).

## Directory Buffers and File Cache Buffers

There are two kinds of cache: directory cache and file cache.

Directory cache is used only with the traditional NetWare file system, not with the Novell Storage Services (NSS) file system. The directory cache is a percentage of the file cache assigned to the traditional file system. Directory cache stores disk directory entries. When the server is first installed, NetWare allocates 20 directory cache buffers of 4 KB each. If more buffers are needed, NetWare tunes the server by increasing the number of directory cache buffers.

You can tune the number of directory cache buffers yourself, using two server parameters that establish the minimum and maximum number of directory cache buffers the server can allocate. For information about tuning directory cache, see [“Tuning Directory Cache for the Traditional File System” on page 50](#).

File cache is used with both the traditional file system and the NSS file system. File cache is a pool of 4 KB memory pages that stores recently used file data.

In an ideal environment, there is enough memory so that all frequently used data can be retrieved from file cache rather than from disk.

Monitoring the disk cache utilization statistic can help you estimate whether your server has sufficient memory to service requests from cache. For information about optimizing file cache, see [“Tuning File Cache” on page 45](#).

If you are using both Novell Storage Services (NSS) and the traditional file system, you want to tune both the file and directory caches. If you are using just the traditional file system, you should tune only the directory cache.

## Garbage Collection

NetWare uses memory deallocation and garbage collection to collect previously allocated but currently unused segments of memory and return them to a common memory pool. Garbage collection is used by the traditional file system only.

Garbage collection is triggered by demand on the virtual memory system. The virtual memory system is checked every 10 seconds. For an explanation of virtual memory, see [“Virtual Memory” on page 14](#).

If virtual memory is being used heavily, garbage collection happens immediately. If virtual memory usage is low, garbage collection happens according to the following process:

1. A process named Free labels the memory as deallocated.
2. The garbage collection process recovers the pieces of deallocated memory and returns them to the memory pool.

The garbage collection is interruptible, can run in the background, and should run frequently. You can manually initiate the process if needed. See [“Optimizing Garbage Collection” on page 55](#).

## Logical Memory Addressing

The NetWare logical memory addressing scheme helps to prevent memory fragmentation. The physical memory of the server is divided into 4 KB blocks called *pages*. When an NLM program wants to store information in memory, it requests memory pages that are contiguous, so that all the information can be stored in a sequential group of pages.

As contiguous blocks of memory pages are allocated and then made free, small segments of memory are often left over. After a while, there are many such isolated pieces of memory too small to be used. This situation is called *memory fragmentation*. When memory is fragmented, the server often cannot satisfy memory requests even when a lot of memory is available.

The memory management subsystem provides a logical space manager that lists logical memory addresses from 0 to 4 GB. The logical space manager lists addresses up to 4 GB even if the server contains only the minimum 256 MB of RAM.

When an NLM program needs a contiguous block of memory, the operating system finds a set of logical, contiguous addresses in the logical space. Because the logical space has so many logical addresses (one million entries—one for each 4 KB page), the operating system can almost always find the contiguous logical memory addresses the NLM program needs. Then the operating system allocates physical memory, which might or might not be contiguous. The operating system then binds the logical addresses to the physical addresses.

## Protected Address Spaces

A portion of the logical address space is set aside for protected address spaces (sometimes called *user address spaces* or *ring 3*). This protected address space is a portion of the logical address space that has carefully controlled interaction with the server operating system. All protected address spaces use virtual memory. For general information about virtual memory, see [“Virtual Memory” on page 14](#).

A key benefit of protected address spaces is that you can use them to take advantage of virtual memory or to run untried or troublesome applications. Because modules loaded into a protected address space can't corrupt the operating system or cause server abends, the protected address space provides a safe place to run applications.

The operating system address space (ring 0) is sometimes called the *OS address space* or the *kernel address space*. The operating system itself cannot run in a protected address space.

The memory protection marshalling layer, in conjunction with the memory protection subsystem, prevents modules in a protected address space from having direct access to anything outside the address space.

In particular, the memory protection marshalling layer serves as the interface between the protected address spaces and the server operating system. These layers prevent NLM programs in protected spaces from making calls or passing parameters to the operating system that would


corrupt or fault the core operating system. Likewise, programs using memory outside the boundary cannot allocate areas of memory inside the boundary.

You can load modules into a protected address space, unload modules from a space, delete a space, or kill a space. See [“Using Protected Address Spaces” on page 38](#) and [“Loading Protected Address Spaces” on page 39](#).

## Protected Address Spaces and Logical Memory Addressing

Each protected address space has its own set of page translations to provide logical memory addressing. The memory management subsystem ensures that the set of page translations for each protected address space refers to a different set of physical memory addresses.

Only one set of translations is visible in the memory management subsystem at a time. When there is more than one protected address space, the memory management subsystem loads one set of page translations for a specified time and then replaces it with another.



Replacing visible page translations is part of a context switch, which takes place when the NetWare kernel changes from running one thread to running another thread.

## Modules Not Allowed in Protected Address Spaces

The following NLM programs and executable files cannot be loaded into a protected space:

- ♦ Server.exe
- ♦ LAN and disk drivers

Hardware and LAN drivers make calls to the operating system and the processor that are not allowed from the protected address spaces.

- ♦ Monitor.exe

Monitor makes calls to the operating system that are not allowed from protected address spaces.

- ♦ Modules that use interrupt time processing
- ♦ Modules that use direct operating system interfaces

- ◆ Modules that import operating system data items
- ◆ Other modules or applications that aren't designed to run in protected address spaces

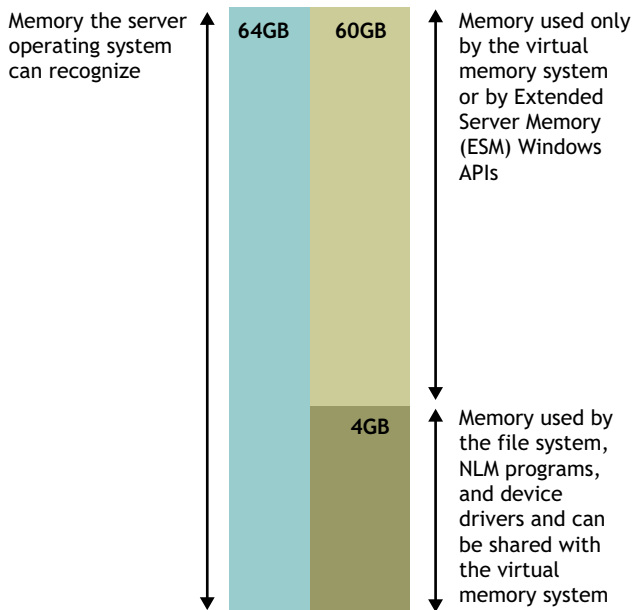
## Virtual Memory

NetWare also provides a virtual memory system that moves little-used data to disk, thereby freeing memory for more frequently used data. Virtual memory provides more efficient use of memory and lessens the likelihood that low memory conditions will cause problems with the server.

Applications that run in protected address spaces use virtual memory, as does Java\*. Applications running in protected space can allocate more memory than if they were not in protected space. Running in protected space keeps an application from corrupting other applications or the kernel; the application is also protected from other applications.

Keep in mind that while virtual memory can be safer for applications and increases the amount of memory that can be used on the server, it also runs slightly slower.

The following illustration shows the amount memory the server can recognize and what portions the virtual memory system uses.



Server operating system modules do not use virtual memory because of the nonpre-emptive nature of the NetWare kernel.

**NOTE:** If an NLM program were designed to use virtual memory and it implemented the needed locking schemes, a NetWare kernel NLM program could be written to use virtual memory.

A periodic aging process assesses available memory to see which data has been used less recently than the rest. Data that has not been used for some time can be moved from memory to disk, thereby freeing memory for other uses.

When the data on disk is needed again, it is moved back into available memory. Because data is swapped in and out of the disk, the available memory can be used for a larger amount of data than its actual physical capacity would allow.

The memory management subsystem assesses the memory needs of all server and NLM processes and determines where the freed memory should be used.

The Java Virtual Machine and any modules that are loaded into protected address spaces use virtual memory. Both the modules and the data they access are subject to being swapped to disk.

When data is moved from memory to disk, it is stored in a swap file. A swap file is created for volume sys: by default.

To create a swap file, delete a swap file, change swap file parameters, or view information about the swap file, you can use Novell Remote Manager or the swap command at the System Console prompt. See [“Using Virtual Memory” on page 52](#).

Swapping the data between memory and disk requires additional system resources or CPU time, but it increases the memory available for use. The availability of the additional memory allows more NLM programs to run in the same amount of memory but they will probably run more slowly. Overall server performance might improve, but probably won't.

## What's Next

For information on managing server memory, begin with [“Using Server Memory Management Tools” on page 19](#).

To increase the efficiency of server memory, see [Chapter 4, “Optimizing Server Memory,” on page 43](#).





# 2 What's New

This section includes the features that were updated in the NetWare® server memory system since its release in NetWare 6.0:

- ♦ [NetWare 6.5 Support Pack 3 Release \(page 17\)](#)
- ♦ [NetWare 6.5 Support Pack 2 Release \(page 17\)](#)
- ♦ [NetWare 6.5 Initial Release \(page 17\)](#)

## NetWare 6.5 Support Pack 3 Release

NetWare 6.5 Support Pack 3 includes the following improvements for server memory:

- ♦ The Auto Tune Server Memory SET parameter was added to automatically tune server memory. This parameter is hidden and turned on by default.
- ♦ On the View Memory Config page in Novell® Remote Manager, the graphs were combined and changed to display more visibly how server memory is used in NetWare. Some link names in the table were simplified.

## NetWare 6.5 Support Pack 2 Release

NetWare 6.5 Support Pack 2 includes the following improvements for server memory:

- ♦ Tools in Novell Remote Manager for monitoring Logical Address space health and use trends (see [“Health Monitor Page” on page 20.](#))
- ♦ Tool in Novell Remote Manager for tuning Logical Address Space usage (see [“Tuning Logical Address Spaces” on page 49.](#))

## NetWare 6.5 Initial Release

NetWare 6.5 includes the following improvements for server memory:

- ♦ Memory moves quicker from one cache pool to another
- ♦ Improved Novell Remote Manager tools for reviewing memory

For more information, see [“Viewing Memory Configurations”](#) in the *NetWare 6.5 Novell Remote Manager Administration Guide for NetWare*.



# 3

## Managing Server Memory

This section briefly describes the tools you use to manage memory and how to protect your operating system by using protected address spaces.

- ◆ [Using Server Memory Management Tools \(page 19\)](#)
- ◆ [Using Protected Address Spaces \(page 38\)](#)
- ◆ [Loading Protected Address Spaces \(page 39\)](#)
- ◆ [Unloading Protected Address Spaces \(page 41\)](#)
- ◆ [Viewing Information about Address Spaces \(page 42\)](#)

### Using Server Memory Management Tools

NetWare® provides the following tools to help you manage and tune the memory system:

- ◆ [Using Novell Remote Manager \(page 19\)](#)
- ◆ [Using Set Parameters \(page 38\)](#)
- ◆ [Using Command Line Commands \(page 38\)](#)

### Using Novell Remote Manager

Novell® Remote Manager contains several pages that can help you understand how your server is using physical and virtual memory and which programs are using it. It can also help you to understand how the logical address spaces for the NetWare kernel are being consumed. Knowing how memory is used can help you determine whether to add more memory to your server to increase server performance, to unload specific NLM™ programs that might be causing memory problems, or to change other settings that might allow the server to use memory more proficiently.

You can also use Novell Remote Manager to load and unload modules to/from protected address spaces, change SET parameter values, and access the server console to execute any console commands.

Novell Remote Manager is more useful than Monitor when diagnosing server memory problems. For more information, see the *NetWare 6.5 Novell Remote Manager Administration Guide for NetWare*.

This section discusses the following pages in Novell Remote Manager:

- ◆ [Health Monitor Page \(page 20\)](#)
- ◆ [List Modules Page \(page 21\)](#)
- ◆ [View Memory Config Page \(page 22\)](#)
  - ◆ [Virtual Memory Health Information Page \(page 30\)](#)

- ◆ [Swap File Information Page \(page 31\)](#)
- ◆ [Virtual Memory Cache Pool Information Page \(page 34\)](#)
- ◆ [Virtual Memory Statistical Information Page \(page 35\)](#)

## Health Monitor Page

### Purpose

Viewing the Health Monitor page in Novell Remote Manager can give you an immediate view of key indicators for monitoring server memory. You can also set thresholds for the monitoring these items and view trend graphs for these items.

### Access

To access this page, click Health Monitor under the Diagnose Server heading in the navigation frame.

### Information Available

**Health Monitor** ?

Begin Refresh Page Refresh Rate 10 seconds

[Mail Notification Configuration](#)
[Threshold Configuration](#)
[Statistic Trend Graphs](#)

Server Health Table							
Include	Notify	Info	Description	Status	Current	Peak	Max
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Work To Do Response Time</a>		0	0	0
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Allocated Server Processes</a>		10	10 *	750
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Available Server Processes</a>		9	7 *	750
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Abended Thread Count</a>		0	N/A	N/A
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">CPU Utilization-0</a>		1	100	100
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Failed Logins Per Hour</a>		0	0	N/A
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Connection Usage</a>		15	17	N/A
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Available Memory</a>		319MB	248MB *	511MB
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Available Logical Space</a>		2,507MB	2,497MB *	3,072MB
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Virtual Memory Performance</a>		N/A	N/A	N/A
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Cache Performance</a>		N/A	N/A	N/A
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">DS Status</a>		LOADED/OPEN	N/A	N/A
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">DS Thread Usage</a>		0	0	375
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Packet Receive Buffers</a>		2000	2000 *	10000
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Available ECB's</a>		1974	1963 *	10000
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">LAN Traffic</a>		6	2742	N/A
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Vol Sys Available Disk Space</a>		2,022MB	1,934MB *	3,934MB
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Available Disk Space</a>		2,759MB	2,759MB *	3,170MB
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Available Directory Entries</a>		14,584	14,584 *	30,968
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Disk Throughput</a>		2	3144	N/A
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Support Pack Status</a>		Not Needed	N/A	N/A
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Timesync Status</a>		LOADED/IN SYNC	N/A	N/A
<input checked="" type="checkbox"/>	<input type="checkbox"/>		<a href="#">Licensing Information</a>		3	3	N/A

\*Low Value

The following key indicators for monitoring server memory are available.

- ◆ Available Memory
- ◆ Available Logical Space
- ◆ Virtual Memory Performance
- ◆ Cache Performance

You can also set thresholds for monitoring these items and view trend graphs for available memory and available logical space.

To change a thresholds from the default values:

- 1 Click the Threshold Configuration link near the top of the page.
- 2 Click the applicable threshold link.
- 3 In the applicable fields that you want to change, type the new values.
- 4 Click OK.

To view a trend graph:

- 1 Click the Statistic Trend Graphs link near the top of the page.
- 2 Select the applicable check boxes for the trends that you want to view.
- 3 Click Draw Selected Graphs.

## List Modules Page

### Purpose

The primary purpose of this page is to view a list of the NetWare Loadable Module™ (NLM) programs loaded on the server. However, the default sort for this page shows which programs are using the most total memory and how that memory is allocated and used by cache, stack, code, and data. Therefore, excluding nss.nlm and server.nlm, which always consume the most memory, you can quickly see which programs are using the most memory.

If the allocated memory keeps growing for an NLM, that program probably has a memory leak and you should quit running it as soon as possible.

### Access

To access this page, click List Modules link under the Manage Applications heading in the navigation frame.

### Information Available

**Diagnose Server**

- [Health Monitor](#)
- [Profile / Debug](#)
- [Reports / Log Files](#)

**Manage Server**

- [Volumes](#)
- [Partition Disks](#)
- [Console Screens](#)
- [Connections](#)
- [Set Parameters](#)
- [Schedule Tasks](#)
- [Console Commands](#)
- [View Memory Config](#)
- [View Statistics](#)
- [Down / Restart](#)
- [Managed Server List](#)
- [Licensing Information](#)
- [Configure TCPIP](#)
- [IP Address](#)
- [Management](#)

**Manage Applications**

- [List Modules](#)
- [Protected Memory](#)
- [System Resources](#)
- [NetWare Registry](#)
- [Protocol Information](#)

**List Modules** ?

Begin Refresh  10 seconds

[Module Load Search Paths](#)

Load Module   Display Logger Console Screen for Module Load

Module Listing									
Name ▼	Ver	Date ▼	NLM Total ▼	Alloc ▼	Cache ▼	Stack ▼	Code ▼	Data ▼	
<a href="#">SERVER.NLM</a>	5.70.02	May 10, 2004	4,347,262,402	<a href="#">49,769,856</a>	4,293,172,608	1,061,560	1,394,538	1,863,840	
<a href="#">NSS.NLM</a>	3.22	May 10, 2004	275,098,704	<a href="#">22,102,016</a>	252,207,104	0	182,240	595,056	
<a href="#">DS.NLM*</a>	10551.03	Apr 13, 2004	8,758,144	<a href="#">5,734,400</a>	0	64,768	2,460,932	485,756	
<a href="#">CIOSSHIM.NLM</a>	1.11	Mar 26, 2004	7,666,899	<a href="#">5,726,208</a>	1,904,640	0	13,287	22,764	
<a href="#">UIMON.NLM*</a>	1.02.01	May 7, 2004	5,867,064	<a href="#">5,439,488</a>	0	327,680	30,752	69,144	
<a href="#">JAVA.NLM</a>	1.42.02	May 7, 2004	5,778,973	<a href="#">454,656</a>	0	5,029,888	211,221	50,440	
<a href="#">NWUTIL.NLM</a>	1.03.05	Nov 21, 2003	3,417,580	0	3,235,840	0	75,992	105,748	
<a href="#">NWMKDE.NLM</a>	7.94	Dec 11, 2001	3,179,737	<a href="#">2,662,400</a>	0	102,400	343,381	63,364	
<a href="#">PORTAL.NLM</a>	3.00	May 10, 2004	3,159,249	<a href="#">1,482,752</a>	0	98,280	924,857	395,312	
<a href="#">JVM.NLM</a>	1.42.02	May 7, 2004	2,613,855	<a href="#">36,864</a>	0	0	2,182,351	394,640	

In terms of memory, this page shows or gives you access to the following information:

<b>Memory Type</b>	<b>Description</b>
Alloc	How much memory an NLM has requested from the system and how much it is currently using. The link takes you to a detailed page that shows the breakdown of how the memory is used.
Cache	How much memory is allocated directly from the file system to cache for each NLM.
Stack	How much memory is allocated to the NetWare stack for each NLM.
Extended	How much extended memory an NLM is using.
Code	How much of the allocated memory is dedicated to the code for each NLM.
Data	How much of the allocated memory is dedicated to data for each NLM.

The allocated memory values are helpful in determining whether the module is using too much system memory.

The other values help you to see how much memory each NLM needs in order to run.

## **View Memory Config Page**

### **Purpose**

This page is helpful when you want to get an overall idea of how the physical and logical memory is being used in your server.

### **Access**

To access this page, click View Memory Config under the Manage Server heading in the navigation frame.


## Information Available

### View Memory Config

Physical Server Memory		
Memory Usage	Actual	Peak Value
<a href="#">NLM Memory</a>	1,817,554,944	1,829,474,304
<a href="#">File System Cache</a>	804,130,816	862,240,768
Extended File System Memory	1,715,470,336	1,715,470,336
<a href="#">Virtual Memory System</a>	20,910,080	25,571,328
Extended Virtual Memory System	47,357,952	47,357,952
Available Cache Memory	1,110,904,832	1,086,943,232*
DOS, Video, Misc. Memory	4,166,656	4,166,656
<a href="#">Extended Server Memory In Use</a>	1,887,436,800	1,887,436,800
<a href="#">Total System Memory</a>	12,347,597,824	N/A

Logical Address Space		
Space Usage	Actual	Percent
In Use Kernel Space	2,831,511,552	65.9
Available Kernel Space	141,787,136	3.3
Fragmented Kernel Space	113,709,056	2.6
Reserved for User Spaces	1,040,187,392	24.2
Reserved for OS Data	167,772,160	3.9
Peak User Space	822,902,784	19.1
Total Address Space	4,294,967,296	100.0


[Tune Logical Address Space](#)



The following table describes the statistic or action of each link shown on the View Memory Config page.

Statistics or Link	Description or Action
NLM Memory	<p>Shows the total physical memory allocated by the operating system to the NLM programs that are loaded on the server.</p> <p>Links to the List Modules page that shows the details of how each NLM program is using memory.</p>
File System Cache	<p>Shows the amount of physical server memory in use by the directory cache buffers in the NSS or traditional NetWare file system.</p> <p>Links to the File System Cache Statistics page that shows the details of these cache statistics. Monitoring these statistics over time helps you determine when you might need to install more RAM for cache. See the online help for details of each of the statistics on this page.</p>
Extended File System Memory	Shows the amount of physical extended memory (above 4 GB) that is in use by the NSS file system.
Virtual Memory System	<p>Shows how much physical memory the virtual memory system is using. The state of the virtual memory system can quickly be determined by the amount of physical memory backing the virtual memory system.</p> <p>Links to the Virtual Memory Health Information page:</p> <ul style="list-style-type: none"> <li>◆ <b>Physical memory:</b> The amount of physical memory backing the virtual memory system.</li> <li>◆ <b>Memory swapped to disk:</b> The amount of virtual memory swapped to disk and not backed by the virtual memory system.</li> <li>◆ <b>Total bytes in virtual memory:</b> The amount of virtual memory that has been used.</li> </ul> <p>Link to Other Pages:</p> <ul style="list-style-type: none"> <li>◆ <b>Swap File Information page:</b> Shows the amount of hard disk space that has been allocated as a virtual memory swap file.</li> </ul> <p>If you find that your server is swapping to virtual memory and its performance is suffering because of it, add more memory to the server.</p> <ul style="list-style-type: none"> <li>◆ <b>Virtual Memory Cache Pool page:</b> Shows statistics for the physical memory cache pool, virtual memory cache pool, and individual cache pools. See the online help on this page for details about each of these statistics.</li> <li>◆ <b>Virtual Memory Statistics page:</b> Shows specific statistics related to virtual memory and provides links to the Virtual Memory Cache Pool page and Virtual Memory Health page.</li> </ul>
Extended Virtual Memory System	Shows the amount of extended (above 4 GB) physical memory that the virtual memory system is using.
Available Cache Memory	Shows the amount of cache memory the server has available.
DOS, Video, Misc. Memory	Shows how much physical memory is in use before the NetWare operating system loads.



Statistics or Link	Description or Action									
Extended Server Memory In Use	In servers that have greater than 4 GB of physical memory loaded, this shows the amount of memory in use and links to the Extended Server Memory Information page that shows the amount of memory in use and available amounts above 4 GB, along with a usage chart.									
Total System Memory	Shows how much total physical memory the server has.									
Logical Address Space Usage table	Shows the logical address space information for the NetWare kernel including space in use, available, fragmented, and reserved for User Spaces and OS data. It also shows the peak usage for User Space and the total amount of address space. The User Space reserve is commonly known as Protected Address Space and it always reserves 1 GB of memory whether it is used or not.									
Tune Logical Address Space link	<p>Displays the Tune Logical Address Space page that shows the Set parameters to change for best memory performance under the current conditions for that specific server.</p> <p>To make the changes, click the Change link next to the parameter you want to accept the changes for and restart the server with the -u command displayed.</p> <p><b>Tune Logical Address Space</b> </p> <table border="1"> <thead> <tr> <th colspan="3">Recommended Set Parameter Changes</th> </tr> <tr> <th>Set Parameter</th> <th>Current Value</th> <th>New Value</th> </tr> </thead> <tbody> <tr> <td>File Cache Maximum Size</td> <td>1,250,000,000</td> <td>1,187,000,000 <a href="#">Change</a></td> </tr> </tbody> </table> <p>Restart the server after making the changes using the following command:</p> <pre>server -u939000000</pre> <p>If you want this change to be permanent, insert this command in the appropriate configuration file.</p> <p>Caution: The recommended changes are for the current server only. After changing parameters, you need to let the server run for a while to adjust to the services being run with the new settings.</p> <p>Use caution when changing these parameters because they are determined to provide the best results for the server while it has been running. If the way the server is used changes, you might have to reset these settings.</p> <p>The settings are generated dynamically for the specific service at the use level of the server and should not be replicated to any other server.</p>	Recommended Set Parameter Changes			Set Parameter	Current Value	New Value	File Cache Maximum Size	1,250,000,000	1,187,000,000 <a href="#">Change</a>
Recommended Set Parameter Changes										
Set Parameter	Current Value	New Value								
File Cache Maximum Size	1,250,000,000	1,187,000,000 <a href="#">Change</a>								

Statistics or Link	Description or Action
Bar Chart	<p>Shows how the logical address space is mapping the physical memory on this machine.</p> <ul style="list-style-type: none"> <li>◆ Memory Map: Shows how the logical address space is set up to run on this machine.</li> <li>◆ Logical Space: Shows how the actual logical address space is being used on this machine.</li> <li>◆ Physical Memory: Shows how the physical memory with addresses smaller than 4 GB are being used on this machine.</li> <li>◆ Extended Memory: Shows how the physical memory with addresses larger than or equal to 4 GB are being used on this machine.</li> <li>◆ L=P: The amount of logical address space that is reserved for mappings of logical equal physical memory. This memory is primarily used for file cache, but some drivers and NLM programs still require it. This size of the logical equals physical memory pool can be changed by using the File Cache Maximum Size Set parameter. Increasing this pool decreases the NLM Memory pool.</li> <li>◆ FS Cache: The amount of memory being used for the file system cache. This color on the Extended Memory bar chart shows the amount of extended memory being used by the file system.</li> <li>◆ Frag: The amount of logical address space that is still available but is in sizes smaller than 4 MB.</li> <li>◆ Avail: The amount of logical address space or physical memory that is available to be used by the system.</li> <li>◆ NLM: The amount of logical address space or physical memory that is being used by NLM programs.</li> <li>◆ Peak US: The peak amount of logical address space that has been used in the User address space. On the Physical Memory and Extended Memory chart, this is the amount of physical memory backing all applications running in User space.</li> <li>◆ US: The amount of logical user address space reserved for running applications in User space or ring 3. You can change the size of the User address space by using the -u switch when issuing the server command. For example, to set the User space size to 500 MB, use the following command when starting the server: <code>server -u500000000</code></li> <li>◆ OS Data: The amount of logical address space reserved on the Logical Space bar or physical memory being used on the Physical Memory bar for various Operating system data. This data is used to manage both the physical and virtual memory being used on the system.</li> </ul>

## File System Cache Statistics Page

### Purpose

This page lets you view File System Cache statistics for this NetWare server. Cache utilization statistics can help you determine when you need to install more RAM for cache.

### Access

To access this page in, click View Memory Config under the Manage Server heading in the navigation frame, and then click File System Cache.

## Information Available

### File System Cache Statistics

Begin Refresh

Page Refresh Rate

#### NSS Cache Statistics

Cache Balance Percent	85 %
Cache Hits	2,084,461
Cache Misses	46,258
Cache Hit Percentage	98 %
Cache Buffers	512
Cache Pages Allocated	450,085

#### Traditional File System Cache Statistics

Short Term Cache Hits	100 %
Short Term Cache Dirty Hits	100 %
Long Term Cache Hits	94 %
Long Term Cache Dirty Hits	0 %
LRU Sitting Time	01:23:54:24:7
Allocate Block Count	53
Allocated from AVAIL	53
Allocated from LRU	0
Allocate Wait	0
Allocate Still Waiting	0
Too many Dirty Blocks	0
Cache ReCheck Block Count	0

The following table describes the cache statistics for Novell Storage Services™ (NSS).

<b>NSS Cache Statistic</b>	<b>Description</b>
Cache Balance Percent	<p>The percent of cache between NSS and the server.</p> <p>The cache balance percentage determines how many cache blocks NSS will take from the traditional file system for its own cache.</p> <p>Every 30 seconds (by default) or when the server loads an NLM, NSS rebalances the NSS file system cache to the percentage specified by the <code>nss /CacheBalance</code> parameter.</p> <p>You can change the default time interval with the <code>nss /CacheBalanceTimer</code> command. Range: 1 to 3,600</p> <p>A high cache balance percentage will impede the performance of the traditional file system. A low cache balance will impede the performance of NSS.</p> <p>We recommend that you set the cache balance parameter to equal the percentage of the total disk space you allocate for NSS. However, do not exceed 80 percent. You might want to initially set the parameter at 50 percent and adjust this number as necessary.</p> <p>Default: 60, Range: 1 to 99</p>
Cache Hits	<p>Number of requests in the last second for disk blocks that were already in cache memory.</p> <p>If you allocate more memory to NSS, the cache hit percentage increases. You can add more memory by adding more physical memory, increasing the cache balance, or increasing the minimum number of cache buffers for NSS. We do not recommend lowering the minimum number of OS cache buffers.</p>
Cache Misses	<p>Number of requests in the last second for disk blocks that were not in cache memory.</p>
Cache Hit Percentage	<p>Percentage of requests in the last second for disk blocks that were in cache memory.</p> <p>This is the number of cache hits divided by the number of attempts to get a file cache buffer from the cache. Any amount above 90 percent is good. This means the system is spending less time going to the disk to retrieve data. Lower numbers indicate NSS needs more memory. If this occurs, you might want to set the cache balance percentage higher.</p>
Cache Buffers	<p>Number of cache buffers used by NSS.</p> <p>If you allocate too much memory to your NSS cache, you might run the server out of memory. To prevent this, set a minimum amount of RAM aside for the exclusive use of the server operation system.</p> <p>Default:256, Range: 256 to 1048576</p>
Cache Pages Allocated	<p>Number of pages allocated for the cache.</p>

The following table describes the cache statistics for the traditional file system.

<b>Traditional File System Cache Statistic</b>	<b>Description</b>
Short Term Cache Hits	Percentage of requests in the last second for disk blocks that were already in cache memory. (When the requested data is already in memory, disk reads don't need to be made.) If this value falls below 98 percent, consider installing more RAM for cache. See <i>also</i> compare Long Term Cache Hits.
Short Term Dirty Cache Hits	Percentage of requests in the last second for disk blocks that were already in cache memory but were dirty. (Dirty cache must be written to disk before being used.) See <i>also</i> Long Term Dirty Cache Hits and LRU Sitting Time.
Long Term Cache Hits	Cumulative percentage of requests for disk blocks that were already in cache. (When the requested data is already in memory, disk reads don't need to be made.) Use the cumulative percentage to assess overall disk cache utilization. If this value falls below 90 percent, install more RAM for cache. See <i>also</i> LRU Sitting Time.
Long Term Dirty Cache Hits:	Cumulative percentage of requests for disk blocks that were already in cache memory but were dirty. (Dirty cache must be written to disk before being used.) Use the cumulative percentage to assess overall disk cache utilization. If this value is high or steadily incrementing, add more RAM for cache. See <i>also</i> LRU Sitting Time.
LRU Sitting Time	<p>Time elapsed since the block that has been sitting longest in the LRU (Least Recently Used) list was referenced, displayed as DD:HH:MM:SS:TS (days:hours:minutes:seconds:tenths of seconds). Long periods of time indicate that either the server has enough memory to handle requests or the request rate is fairly constant. Short periods of time indicate the server might need more RAM for cache.</p> <p>If the server has been up for over 24 hours, this statistic should never drop below 24 hours. It is important to understand that LRU Sitting Time is a measure of noncache memory, not file system cache memory.</p>
Allocate Block Count	<p>Cumulative number of requests for disk cache blocks that have been made since the server was started or rebooted. This value is the sum of the values of Allocated from AVAIL and Allocated from LRU.</p> <p>If the value of Allocated from AVAIL is much higher, the server has sufficient RAM for cache. If the value of Allocated from LRU is high, install more RAM for cache.</p>
Allocated from AVAIL	<p>Number of requests for disk cache blocks that were filled by blocks in the available list (blocks that were not being used).</p> <p>When there are no free blocks available, requests are filled from the Least Recently Used (LRU) list of cache blocks. If this value is much higher than the Allocated from LRU value, the server has sufficient RAM for cache.</p>
Allocated from LRU	<p>Number of requests for disk cache blocks that were filled by blocks from the Least Recently Used (LRU) cache blocks.</p> <p>The system writes pending write requests from the LRU cache block to disk and then frees the block for the current request. Because LRU cache is used only when no other cache is available, a steadily incrementing count indicates more RAM is needed.</p>

<b>Traditional File System Cache Statistic</b>	<b>Description</b>
Allocate Wait	<p>Number of times a request waited while an LRU block was freed.</p> <p>Because LRU cache is used only when no other cache is available, a high or steadily incrementing count indicates that not enough cache memory is available. Install more RAM for cache.</p>
Allocate Still Waiting	<p>Number of times in the last 10 minutes that the operating system waited for an LRU block in order to fulfill a request.</p> <p>If this value is greater than 7, install more RAM for cache.</p>
Too Many Dirty Blocks	<p>Number of times a write request was delayed because there were too many writes to perform or the disk channel was busy.</p> <p>A high value indicates that either the disk channel has too much I/O traffic or you need to install more RAM for cache.</p>
Cache ReCheckBlock Count	<p>Number of times a disk cache request had to be retired because the target block was being used.</p> <p>If this value is high or steadily incrementing, install more RAM for cache.</p>

## Virtual Memory Health Information Page

### Purpose

The state of the virtual memory system is determined by the amount of physical memory backing the total virtual memory system.

If your server is running slow, check the swap space and the amount of physical memory that is backing the physical memory. Keep the swap space as low as possible. We recommend having twice as much physical memory as swap space.

If you continue to need more swap space to keep programs running, you probably need to add more physical memory to the server.

### Access

To access this page, click View Memory Config under the Manage Server heading in the navigation frame, and then click Virtual Memory System.

## Information Available

### Virtual Memory Health Information




Begin Refresh Page Refresh Rate 10 seconds

[\[To Swap File Information\]](#)

[\[To Virtual Memory Cache Pool\]](#)

[\[To Virtual Memory Statistics\]](#)

Virtual Memory Usage	
Memory Usage	Actual
Physical memory in use	1,317,191,680
Memory swapped to disk	0
Total bytes in virtual memory	1,317,191,680



The following key statistics are shown on the Virtual Memory Health Information page:

- ◆ **Physical memory in use:** Amount of physical memory backing the virtual memory system.
- ◆ **Memory swapped to disk:** Amount of virtual memory swapped to disk that is not backed by physical memory.
- ◆ **Total bytes in virtual memory:** Amount of virtual memory in use.

The Virtual Memory Health Information page also contains links to additional pages that can help you understand and manage the virtual memory system in the server:

- ◆ The To Swap File Information link goes to the [Swap File Information Page \(page 31\)](#).
- ◆ The To Virtual Memory Cache Pool link goes to the [Virtual Memory Cache Pool Information Page \(page 34\)](#).
- ◆ The To Virtual Memory Statistics link goes to the [Virtual Memory Statistical Information Page \(page 35\)](#).

## Swap File Information Page

### Purpose

This page shows the amount of hard disk space that has been allocated as a virtual memory swap file. Most problems with virtual memory stem from not having enough space designated for the swap file; therefore, make sure you have enough disk space for this.

When the amount of available memory for swapping becomes low, the virtual memory system begins to struggle. If there is no memory available, then the virtual memory system is able to back memory using only physical memory resources. Under these conditions, the virtual memory system will most likely have a performance degradation.

If you find that your server is swapping to disk and its performance is suffering because of it, add more memory to the server.

## Access

To access this page, click View Memory Config under the Manage Server heading in the navigation frame, and then click Virtual Memory System > To Swap File Information.

## Information Available

### Swap File Information

[\[Back to System Memory Information\]](#)

Swap File Information								
Volume Name	Volume Size	Swapping Enabled	Disk Space			Parameters		
			Reserved for Swapping	Used for Swapping	Available for Swapping	Minimum Reserved	Maximum used for Swapping	Minimum Volume Free Space
SYS	4108 M	<a href="#">no</a>			2671 M			
VOL1	6852 M	<a href="#">no</a>			6201 M			
NSS1	6195 M	<a href="#">yes</a>	70 M	26 M	6145 M	<a href="#">2 M</a>	<a href="#">5874 M</a>	<a href="#">5 M</a>

Summary of disk space utilization by the Virtual Memory System	
Number of volumes currently being used for swapping	1 volume(s)
Total disk space currently reserved for swapping	70 million bytes
Total disk space currently being used for swapping	26 million bytes
Total disk space that is available for swapping	6145 million bytes
Number of volumes not currently being used for swapping	2 volumes
Total disk space available on volumes currently not being used for swapping	8872 million bytes
Page In Count	4 pages
Page Out Count	9945 pages
Swap Reserved Count	16872 pages
Swap Page Count	17280 pages
Swap Bit Map Pages	0 pages
Average Page In Rate	0 pages / second
Average Page Out Rate	0 pages / second

Total Disk Usage



The Swap File Information table shows the following information for each volume in the system.

Heading	Description
Volume Name	Name of each volume.
Volume Size	Size of each volume.
Swapping Enabled	Whether swapping is enabled or disabled. Lets you enable or disable the swap file by clicking the Yes or No links.



Heading	Description
Disk Space Reserved for Swapping	Size (in millions of bytes) that have been reserved for swapping and the amount of memory the virtual memory system needs to back any virtual memory allocation for each volume.
Disk Space Used for Swapping	Amount of memory (in millions of bytes) that is being used by the Virtual memory system for each volume. This is the amount of memory that has been actually paged-out to disk.
Disk Space Available for Swapping	Amount of memory (in millions of bytes) that can still be paged-out to the disk for each volume.
Parameters, Minimum Reserved	Minimum size for the swap file.
Parameters, Maximum Used for Swapping	Sets the maximum size for the swap file. Normally, the swap file is allowed to grow to the size of the volume. However, you can limit the size if necessary.
Parameters, Minimum Volume Free Space	Maximum size for the swap file. This parameter allows the swap file to grow; however, it guarantees that a given amount of disk space will be left free.  For example, if this parameter is set to 10, the swap file will grow only if 10 million bytes on the disk will remain free after the growth.

The Summary of Disk Space Utilization by the Virtual Memory System table shows the following:

Heading	Description
Number of volumes currently being used for swapping	Number of volumes that have a currently active swap file. Each volume can have one or no swap file.
Total disk space currently being used for swapping	How much disk space the virtual memory system is actually using. The number should be less than or equal to the reserved amount.
Total disk space that is available for swapping	How much disk space you can use for all of the swap files.
Number of volumes not currently being used for swapping	Number of volumes that do not have a swap file on them. For performance reasons, it is better if you allow a swap file on all available volumes unless you have a specific reason for not doing so.
Total Disk Space Available on Volumes Currently Not Being Used for Swapping	How much memory could be used for swapping if there were swap files on those volumes that do not currently allow swapping.
Page-In Count	Number of pages that have been brought into the virtual memory system from the disk since the server was last restarted.
Page-Out Count	Number of pages that have been sent to the disk by the virtual memory system since the server was last restarted.
Swap Reserved Count	Number of pages of memory that are reserved in all of the swap files.
Swap Page Count	Number of pages of memory that all of the swap files can back.

Heading	Description
Swap Bit Map Pages	Number of pages of memory the swap file manager is using to manage the bitmaps for swap file control.
Average Page-In Rate	Running 5-second average of the number of pages that the virtual memory system has paged-in from the disk.
Average Page-Out Rate	Running 5-second average of the number of pages that the virtual memory system has paged-out from the disk.

## Virtual Memory Cache Pool Information Page

### Purpose

This page shows statistics for the physical memory cache pool, virtual memory cache pool, and individual cache pools. See the online help on this page for details about each of these statistics. This page is more useful for developers who are trying to understand how memory in the system is being used by an NLM program that they are developing.

### Access

To access this page, click View Memory Config under the Manage Server heading in the navigation frame, and then click Virtual Memory System > To Virtual Memory Cache Pool.

### Information Available

#### Virtual Memory Cache Pool Information

Begin Refresh Page Refresh Rate 10 seconds

[\[To Swap File Information\]](#)  
[\[To Virtual Memory Health\]](#)  
[\[To Virtual Memory Statistics\]](#)

Cache Pool Physical Memory Statistics		
	Count	Percentage
Physical Page Breakdown	<a href="#">Composite Chart</a>	<a href="#">Composite Chart</a>
VM	235,579	44%
FS Cache	292,234	55%
Total Pages In Use	527,813	

VM Cache Pool Specific Statistics	
Pages Transferred	<a href="#">Meter</a>
Last Second	0
VM Cache Pool Free Lists	
Used Pages (not on a free list)	<a href="#">77,781</a>
Free Clean Pages	<a href="#">382,721</a>
Free Cache Pages	<a href="#">0</a>
Free Dirty Pages	<a href="#">0</a>

Individual Cache Pool Statistics		
Name	VM	FS Cache
Type	Virtual Memory	Classic File System
Pages Allocated	<a href="#">460,502</a>	257,532
Average Pages	460,502	234,246
Cache Available State	Extra-Plenty	Extra-Plenty
Average Cache Available State	Extra-Plenty	Extra-Plenty
Transfer Operation	N/A	No Transfer
Pages To Transfer	0	
Control Thread	N/A	<a href="#">0x48245420</a>

## Virtual Memory Statistical Information Page

### Purpose

This page contains virtual memory statistical information. This information shows you what is happening in the virtual memory system when the VM Cache Pool Percentage parameter is set. We recommend setting this parameter and letting it control the system rather than trying to tweak the system yourself.

### Access

To access this page, click View Memory Config under the Manage Server heading in the navigation frame, and then click Virtual Memory System > To Virtual Memory Statistics.

### Information Available

#### Virtual Memory Statistical Information

Page Refresh Rate

[\[To Virtual Memory Cache Pool\]](#)

[\[To Virtual Memory Health\]](#)

Statistical Information	
Current Available Memory	232,193
Free Desired Pages	64
Free Lots Pages	80
Free Minimum Pages	48
Memory Available State	Extra-Plenty
Need Free Clean Pages	0
Need Free Pages	1
Upper Page I/O Threshold	32

The page shows the following virtual memory statistical counters:

Statistic	Description
Current Available Memory	Amount of physical memory that is available to the virtual memory system.
Free Desired Pages	Number of pages that the virtual memory system needs for proper operation.
Free Lots Pages	Number of available pages that the virtual memory system finds excessive. Having an available page count beyond this number can result in pages being given to another cache pool.
Free Minimum Pages	Minimum number of pages that the virtual memory system needs for proper operation.

Statistic	Description
Memory Available State	<p>Current state of the virtual memory system. The six possible states are the following:</p> <ul style="list-style-type: none"> <li>◆ Desperate</li> <li>◆ Poor</li> <li>◆ Fair</li> <li>◆ Normal</li> <li>◆ Plenty</li> <li>◆ Extra-plenty</li> </ul>
Need Free Clean Pages	Flag which, if the value is not 0, indicates that the virtual memory system needs some clean pages to operate properly.
Need Free Pages	Flag which, if the value is not 0, indicates that the virtual memory system needs some free pages.
Upper Page I/O Threshold	<p>Controls an alert that triggers when the server has had too many page I/Os recently.</p> <p>This is an indication that the virtual memory system is thrashing because of too little physical memory resources.</p>

## Server Memory Pool Summary Page

### Purpose

This page shows a summary all of the server memory detailed statistics.

### Access

To access this page, click View Memory Config under the Manage Server Heading in the navigation frame, then click Total System Memory.

## Information Available

### Server Memory Pools Summary ?

Begin Refresh Page Refresh Rate 10 seconds v

Memory Summary Table							
Known Memory	Bytes	Pages	Memory	Kernel NLM Memory	Bytes	Pages	Memory
Server	535,090,368	130,637	510 MB	Code	32,755,712	7,997	31 MB
DOS	75,584	18	73 KB	Data	18,481,152	4,512	17 MB
Video	8,192	2	8 KB	Shared Code	40,960	10	40 KB
Other	131,072	32	128 KB	Shared Data	20,480	5	20 KB
Virtual Memory System	Bytes	Pages	Memory	Messages	1,191,936	291	1 MB
Free Cache	0	0	0 KB	Help	0	0	0 KB
Free Clean	14,209,024	3,469	13 MB	Alloc LI=P	143,372,288	35,003	136 MB
Free Dirty	0	0	0 KB	Alloc L=P	512,000	125	500 KB
In Use	3,686,400	900	3 MB	NLM Total Memory	196,374,528	47,943	187 MB
Virtual Memory Total	17,895,424	4,369	17 MB	Cache Information	Bytes	Pages	Memory
Other	Bytes	Pages	Memory	Cache Buffer Size	4,096	1	4 KB
File System Memory	56,692,736	13,841	54 MB	Original Cache Buffers	531,087,360	129,660	506 MB
Reserved Swap Memory	4,116,480	1,005	3 MB	Available Cache Buffers	260,112,384	63,504	248 MB
Swap File Size	6,029,312	1,472	5 MB	Dirty Cache Buffers	0	0	0 KB

The Server Memory Pools Summary shows the following information for each type of memory in the server. The amounts consumed are shown in bytes of disk space, pages, and memory used.

Memory Type	Description
Known Memory	Distribution of the known memory among DOS, video buffers, and the server memory pool.
Kernel NLM Memory	<p>How NLM programs loaded in the OS address space of the NetWare server are using memory.</p> <p>The Alloc LI=P category shows NLM memory allocations using mapped memory where the logical address of the memory does not equal the physical address of the memory. This is the normal state of most NLM memory allocations.</p> <p>The Alloc L=P category shows memory allocations where the logical and physical addresses of the memory is the same. This type of memory allocation is done mainly for Direct Memory Access (DMA) used by device drivers that cannot handle different logical and physical addresses.</p>
Virtual Memory System	Memory usage by the NetWare server virtual memory system.
Cache Information	<p>Original, current, and dirty cache buffer information.</p> <p>Original Cache Buffers: Amount of memory available to the server cache when the server first came up.</p> <p>Current Cache Buffers: Amount of memory available now for file caching and memory allocations.</p> <p>Dirty Cache Buffers: Server cache memory pages that have data in them that needs to be flushed to disk.</p>

Memory Type	Description
Other	<p>File System memory, reserved swap memory, and the swap file size information.</p> <p>File System Memory: Amount of server memory in use by the directory cache buffers in the traditional NetWare file system.</p> <p>Reserved Swap Memory: Small footprint of reserved memory that allows the NetWare operating system to point to the physical location of the actual swap file.</p> <p>Swap File Size: Amount of hard disk space that has been allocated as a virtual memory swap file.</p>

## Using Set Parameters

Settable parameters and console commands let you tune the memory system to achieve optimal performance. For more information, see SET > “[Memory Parameters](#)” in the *NetWare 6.5 Utilities Reference*.

## Using Command Line Commands

For NSS, you can also enter commands at the System Console.

## Using Protected Address Spaces

You can use a portion of server memory called a *protected address space* to protect the server operating system from untried or troublesome applications and to use virtual memory. When applications are loaded into a protected address space, they can't corrupt the operating system, cause server abends, or corrupt other applications that are running in other protected address spaces.

All protected address spaces use virtual memory, so running modules in a protected address space also uses RAM more efficiently.

For general information about protected address spaces and the applications that can be loaded into them, see “[Protected Address Spaces](#)” on page 12. For general information about virtual memory, see “[Virtual Memory](#)” on page 14.

The following are general guidelines for using protected address spaces:

- ◆ Because modules loaded into a protected address space have controlled interaction with the operating system, all modules that must communicate with each other by directly importing another module's data or functions (such as those of an e-mail application or a database) should be loaded into the same protected address space.
- ◆ If you load a NetWare Loadable Module (NLM) that depends upon another module, such as CLIB, the other module will be loaded automatically into the same address space.
- ◆ If the same module is loaded into more than one address space, the module's code will be shared among the address spaces. Therefore, loading the module into multiple address spaces does not require additional memory for the module itself. Only data for the required module is unique for each address space.
- ◆ Even if an NLM is designed to be loaded only once, you can load multiple copies of the NLM on your server if you load them into different address spaces.

- ◆ When you load modules into a protected address space, NetWare assigns whatever amount of memory the loaded modules need, up to a maximum size of 512 MB. The maximum size of a protected address space is fixed— but within that maximum limit, the memory size of the space grows and shrinks as needed by the modules in the address space.
- ◆ If you want the server to clean up address spaces that abend, in Novell Remote Manager make sure the Memory Protection Fault Cleanup parameter (Memory category) is set to the default setting of On. If an address space abends, the server removes the address space and its modules and returns the resources to the system.

If the Memory Protection Fault Cleanup parameter is set to Off, the situation is left to the NetWare kernel's abend recovery mechanism.

## Loading Protected Address Spaces

You can create a protected address space when you load one or more modules.

**Protected Memory** ?

**Memory Protection Console Interfaces**

Load NCF File Protected

Load Module Protected

Address Space Information							
Address Space Name	State	Time Since Last Fault	Restart Count	Restartable	Unload	Remove	Restart
<a href="#">ADMINSRV</a>	ADDRESS_SPACE_OK	N/A	0	No	Unload	Remove	N/A
<a href="#">Java_Kernel_Space</a>	ADDRESS_SPACE_OK	N/A	0	N/A	N/A	N/A	N/A
<a href="#">Java_Space_343</a>	ADDRESS_SPACE_OK	N/A	0	N/A	N/A	N/A	N/A
<a href="#">Java_Space_454</a>	ADDRESS_SPACE_OK	N/A	0	N/A	N/A	N/A	N/A
<a href="#">Java_Space_477</a>	ADDRESS_SPACE_OK	N/A	0	N/A	N/A	N/A	N/A
<a href="#">Java_Space_703</a>	ADDRESS_SPACE_OK	N/A	0	N/A	N/A	N/A	N/A
<a href="#">OS</a>	ADDRESS_SPACE_OK	N/A	0	N/A	N/A	N/A	N/A

**Memory Protection Set Parameters**

Memory Protection Fault Cleanup ON  
 Memory Protection No Restart Interval 1  
 Memory Protection Restart Count 1  
 Memory Protection Abend After Restart Count OFF  
 Auto Restart After Abend 1

To load protected address spaces, use the commands in the following table at the System Console prompt or in the Memory Protection Console Interface fields on the Protected Memory page in Novell Remote Manager (accessed by clicking Protected Memory under the Manage Applications heading in the navigation frame).

To	Enter this command at the System Console prompt	Do this on the Protected Memory Page of Novell Remote Manager	Example
Load one module into the next available protected address space.	<code>[load] protected module_name</code>	<ol style="list-style-type: none"> <li>1. Type <i>module_name</i> in the Load Module Protected field.</li> <li>2. Click Load Module Protected.</li> </ol>	<p><code>[load] protected database.nlm</code></p> <p>Creates a new address space called <i>address_spacesn</i>, where <i>n</i> is a numerical number, and loads <i>database.nlm</i> into it.</p>

To	Enter this command at the System Console prompt	Do this on the Protected Memory Page of Novell Remote Manager	Example
<p>Load one module into the next available protected address space with restart functionality.</p> <ul style="list-style-type: none"> <li>Restart functionality means that if the protected space abends, the system closes the space, cleans up its resources, restarts the space, and then reloads the module into it.</li> <li>To prevent the server from restarting a memory space that continues to fault, use the Memory Protection No Restart Interval SET parameter (Memory category).</li> </ul>	<pre>[load] restart module_name</pre>	<ol style="list-style-type: none"> <li>Type <b>restart</b> <i>module_name</i> in the Load Module Protected field.</li> <li>Click Load Module Protected.</li> <li>(Optional) If you want to change the restart state to No, click Yes in the Restartable column of the Address Space Information table.</li> </ol>	<pre>restart database.nlm</pre> <p>Creates a new address space called <i>address_space<sub>n</sub></i>, where <i>n</i> is a numerical number, and loads <i>database.nlm</i> into it.</p>
<p>Load one module into a protected space with a user-defined name.</p> <ul style="list-style-type: none"> <li>Use this command when you want to create your own name for the space.</li> <li>Also use this command when you want to load more than one module into the same protected space.</li> <li>Repeat the command for each module you want to load into the space.</li> <li>You can also include the restart option on the same command line.</li> </ul>	<pre>[load] address space = space_name module_name</pre>	<ol style="list-style-type: none"> <li>Type <b>address space =</b> <i>space_name</i> <i>module_name</i> in the Load Module Protected field.</li> <li>Click Load Module Protected.</li> <li>(Optional) If you want to load more than one module in the same space, repeat Steps 1 and 2 for each module you want to load.</li> </ol>	<pre>[load] address space = db_space database.nlm</pre> <p>Creates a new address space called <i>db_space</i> and loads <i>database.nlm</i> into it.</p> <pre>[load] address space = db_space restart test.nlm</pre> <p>Loads <i>test.nlm</i> into the <i>db_space</i> address space and also adds restart functionality to the address space.</p> <pre>[load] address space = db_space test.nlm</pre> <p>Loads the <i>test.nlm</i> in to the same <i>db_space</i> address space that the <i>database.nlm</i> was loaded into.</p>
<p>Load multiple modules at once using an .ncf file.</p> <p>The command creates a protected space with the same name as the .ncf file and executes the .ncf file to load all the modules into the space.</p>	<pre>protect filename</pre>	<ol style="list-style-type: none"> <li>Type <b>NCF_filename</b> in the Load NCF File Protected field.</li> <li>Click Load NCF File Protected.</li> </ol>	<pre>protect load_db</pre> <p>Creates an address space called <i>load_db</i> and executes <i>load_db.ncf</i> to load modules into the protected space.</p>



# Unloading Protected Address Spaces

To unload a module from a protected space or to unload the space itself, use the commands in the following table at the System Console prompt or in the Memory Protection Console Interface fields on the Protected Memory page in Novell Remote Manager (accessed by clicking Protected Memory under the Manage Applications heading in the navigation frame).

To	Enter this command at the System Console prompt	Do this on the Protected Memory Page of Novell Remote Manager	Example
Unload all modules from the designated address space and remove the space.	<code>unload address space = space_name</code>	Click Unload in the Unload column of the Address Space Information table for the address you want to unload the modules from.	<code>unload address space = address_space1</code>  Unloads all the modules from address_space1, closes address_space1, and then returns its resources to the system.
Remove an address space without unloading the modules first.  <b>IMPORTANT:</b> Because you might experience a delay in unloading the address space or cause other modules to malfunction, use this command only if you know the space can't be unloaded with the <code>unload address space</code> command.	<code>unload kill address space = space_name</code>	Click Remove in the Remove column of the Address Space Information table for the address space you want to remove.	<code>unload kill address space=address_space1</code>  Removes address_space1 without unloading modules and returns its resources to the system.

If a module hangs while it is being unloaded with the `unload address space` command, the system waits a specified time and then displays the following error message:

```
Module_name in address_space did not unload yet. Killing the address space is the only way to remove the NLM. Should the address space be killed? Yes/No.
```

If you answer Yes, the system shuts down the address space without unloading modules from it. If you answer No, the system waits a specified time and displays the message again.

You can designate the amount of time before the message appears by setting the Hung Unload Wait Delay (Error Handling category) parameter. The default is 30 seconds.

In some circumstances, the server might not completely clean up an address space that has faulted. This failure is caused by NLM programs that are not written to allow external cancellation.

When the NetWare operating system cannot clean up an address space after it faults, it sends an alert message to the server console. If you execute the protection command, the address space will be listed as in the clean-up failure state.

Because you cannot unload an address space that is in the clean-up failure state, you must shut down and restart the server to recover the remaining resources from the address space.

# Viewing Information about Address Spaces

You can see which modules are loaded in which address spaces.

## In Novell Remote Manager

- 1 Click Protected Memory under the Manage Applications heading in the navigation frame.
- 2 View the address space information on the Protected Memory page.
- 3 To view the modules that are loaded in a specific address space, click the Address Space Name link for the specific address.

## At the System Console Prompt

- 1 To display a list of loaded modules with their address space names, enter

`modules`

Each NLM name displayed is followed by the name of the address space where the NLM resides, the full utility name, version number, date, and other information. The module names are displayed in the order in which they were loaded.

- 2 To display a list of all address spaces, enter

`protection`

Each address space name that is displayed is followed by a list of all modules loaded into that space, with a short description of each module.

## What's Next

To increase the efficiency of server memory usage, see [Chapter 4, "Optimizing Server Memory," on page 43](#).

# 4

## Optimizing Server Memory

This section identifies tasks that you can perform to optimize server memory, thereby increasing network performance. Optimizing server memory includes the following tasks:

- ♦ [Assessing Server RAM \(page 43\)](#)
- ♦ [Ensuring That Devices Can Access Memory \(page 44\)](#)
- ♦ [Determining Which Processes Use Too Much Memory \(page 44\)](#)
- ♦ [Tuning Server Memory \(page 45\)](#)
- ♦ [Using Virtual Memory \(page 52\)](#)
- ♦ [Optimizing Garbage Collection \(page 55\)](#)

### Assessing Server RAM

Insufficient physical memory (RAM) in the server is a common cause of a slow network. To determine whether your server has sufficient physical memory installed, check the amount of server memory currently available for disk cache.

### Traditional File System

To check the amount of server memory currently available for disk cache, you can use NetWare<sup>®</sup> Remote Manager.

- 1** In Novell<sup>®</sup> Remote Manager, click Health Monitor under the Diagnose Server heading in the navigation frame, and then view status of the following:
  - ♦ Available Memory
  - ♦ Available Logical Space
  - ♦ Virtual Memory Performance
- 2** If the status is Suspect or Bad, view the details of the problem on by clicking the link for the applicable threshold.

For example, for Available Memory, click Available Memory on the Health Monitor page. This link takes you to the View Memory Config page where you can access virtually any of the memory information you need.

**TIP:** If you are unable to add more RAM now. We suggest unloading unneeded NLM<sup>™</sup> programs or removing DOS from memory as a short-term fix. Remember that this treats symptoms only and is not a long-term solution.

## Novell Storage Services

To check the amount of server memory used for disk cache, use the following commands directly at the System Console prompt or remotely through Novell Remote Manager:

- ◆ **nss /status**  
Reports on general caching parameters for NSS.
- ◆ **nss /cachestats**  
Displays statistics for cache buffers used by NSS.
- ◆ **nss /resetstats**  
Resets the cache stats. Use this command to trace statistics from day to day, especially during busy hours. Keep in mind that running a backup alters the statistics.
- ◆ **nss /namecachestats**  
Reports on name cache, including filenames and hit rates.

## Ensuring That Devices Can Access Memory

Some older network boards, such as many ISA and MCA devices, cannot access memory above 16 MB. To make sure these devices have enough memory below 16 MB, make sure that the following SET command is in the startup.ncf file:

```
set reserved buffers below 16 meg = 300
```

The default is 300 buffers. If necessary, you can increase the value, but remember that too many reserved buffers can prevent large volumes from mounting. As soon as possible, upgrade the system to a newer board that can access higher memory addresses. For a discussion of network bottlenecks, see “[Matching Performance Capabilities](#)” in the *NetWare 6.5 Server Operating System Administration Guide*.

NetWare automatically registers memory above 16 MB so it can be recognized by EISA and PCI machines. If memory above 16 MB isn't being recognized, you can register memory manually using the register memory command. For more information, see “[Register Memory](#)” in the *NetWare 6.5 Utilities Reference*.

**IMPORTANT:** Use the register memory command only if absolutely necessary. Manually registering memory can cause memory fragmentation problems. Ideally, you should upgrade the system board so that the NetWare automatic memory registration will work.

## Determining Which Processes Use Too Much Memory

To determine which NLM programs are using excessive amounts of memory (have a memory leak), complete the following steps:

- 1** Watch the memory usage trends for any NLM that you might suspect has a problem.
  - 1a** In Novell Remote Manager, click List Modules under the Manage Applications heading in the navigation frame.
  - 1b** If needed, sort the list by allocated memory usage by clicking Alloc Memory in the table.

**TIP:** Server.nlm and nss.nlm modules will always be high memory users.
  - 1c** View the Allocation Summary page for each NLM that is at the top of the list by clicking Alloc Memory for each NLM.

This page contains the Allocation Summary L!P information and links to sort the memory allocation by Resource Tag or size.

Any NLM that has a memory resource that is constantly growing might have a memory leak.

- 2 Watch for Short term memory allocator is out of memory alerts displayed at the System Console screen or Logger screen or in the system log files.

Often one NLM stands out with an unusually large memory resource.

If you suspect an NLM program has a memory leak, you can unload and reload the NLM to solve the problem temporarily.

To solve the problem permanently, contact the vendor of the program and request that they fix the program.

## Tuning Server Memory

NetWare is a self-tuning operating system; however, there are some things you can do to adjust the server's memory subsystem (as well as other aspects of server operation) to fine tune its performance.

Before attempting to tune a new server's memory, let the server operate at full capacity for a week or two to allow NetWare to self-tune. After the server settles into a predictable pattern of daily use, you can begin to fine-tune the server's memory as described in the following sections.

- ◆ [“Tuning File Cache” on page 45](#)
- ◆ [“Tuning Directory Cache for the Traditional File System” on page 50](#)
- ◆ [“Tuning Logical Address Spaces” on page 49](#)

If you are using both Novell Storage Services™ (NSS) and the traditional file system, you should tune both the file and directory caches.

## Tuning File Cache

File caching speeds access to file data. If you want to tune your NetWare server in general, tune the file cache.

### Novell Storage Services File System

In NetWare 6.5, NSS is always loaded on the server in order to run, and cache memory is always allocated. NSS uses memory differently than the traditional file system does, and you must use specific parameters for tuning your server when using NSS.

In cases where you have upgraded a server to NetWare 6.5 and you have no NSS logical volumes, you need to set the following NSS cache values as noted:

```
nss /noCacheBalance
nss /MinBufferCacheSize=2048
```

In NetWare 6.5, all CD and DVD drives are recognized as NSS devices; therefore, you need to set NSS cache parameters to accommodate the usage of these devices.

To view statistics for cache buffers used by NSS, enter one of the following commands at the System Console prompt:

```
nss /cachestats
```

or

```
nss /reset stats
```

With heavy access on an NSS logical volume, the server might incur some performance issues. NSS has its own caching and parameters to be set during the loading of the NSS modules.

### Cache Hit Percentage

From an overall server performance perspective, the most important statistic to look at is cache hit percentage. This indicates how often NSS looked in cache for a file and found it there. To view the cache hit percentage, use the `nss /cachestats` command.

The long-term cache hit percentage should be above 90. A long-term hit percentage below 90 can indicate you have too little RAM on the server, or that one or more of the following SET parameters is set too low.

- ◆ **NSS Cache Balance Percent Set Parameter:** Controls how much memory is available for NSS to use for cache. The percentage setting determines how many cache blocks NSS takes from the traditional file system for its own cache. A high cache balance percentage impedes the performance of the traditional file system. A low cache balance impedes the performance of NSS.

We recommend that you set the Cache Balance Percent Set parameter to equal the percentage of the total disk space you allocate for NSS. The default is 85 percent. If you are using mostly NSS volumes, you can run at a higher percentage, such as 90 percent.

- ◆ **NSS Close File Cache Size Set Parameter:** Dictates how many closed files are kept in cache. The default is 100,000. If applications on a server consistently open and close many files, increase this setting. If the applications use only a few file large files, decrease this setting so more memory can be used for the buffer cache.

NSS adjusts the number of closed files it keeps in cache to balance the needs of the file system with the needs of other processes. NSS does not automatically keep the number of closed files in cache that are indicated by this parameter. This parameter just indicates a maximum amount.

- ◆ **NSS File Flush Timer and NSS Buffer Flush Timer Set Parameters:** Sets the number of seconds that modified files and buffers stay in memory before being written to disk. Increasing these parameters from their defaults helps the caching up to a certain point, but too much increase causes a *throttling* problem with writing out the metadata. This is somewhat dependent on processor speed and disk speed, so you could try adjusting them for your system. If NSS performance is very poor, 75 might be too high and you could try 50 percent. If performance is just below normal, you could try increasing the values a little.

To activate these parameters, you can issue the following command at the startup of NSS, or place it in the `autoexec.ncf` file:

```
nss /cachebalance=80 /fileflushtimer=75 /bufferflushtimer=75 /  
closedfilecachesize=50000
```

NSS caches files greater than 128 KB in the memory above 4 GB. This improves performance for these files and for smaller files because they won't be competing for the same space in the memory below 4 GB.

For more information on setting Cache Buffers for NSS, see [“Setting the Cache Buffers”](#) in the *Novell Storage Services File System Administration Guide for NetWare 6.5*.

## Assessing Changes to NSS Settings

After making changes to NSS settings, enter the following command at the System Console prompt:

```
nss/cachestats
```

Note the statistics and how long the server has been up. After noting the statistics, enter the following command at the System Console prompt:

```
nss/resetstats
```

This command clears the file system cache statistics and starts keeping new statistics. This allows you to see the result of your changes.

Because it takes time for the system to adjust to the loads, you must allow NSS to keep statistics for at least two full business days before you consider the results valid. If the server has been up for more than a few weeks before you reset the statistics, wait at least a week before you consider the results valid.

Do not add RAM over 4 GB to improve cache performance. The Intel\* Architecture\* has no DMA capabilities above 4 GB; therefore, NSS deliberately does not use any memory above this amount.

## Traditional File System

Because it is much faster to get data from memory than to get it from disk, caching frequently requested data increases server performance. The more RAM installed in the server, the more cache memory is available. Cache hits occur when a client requests data and the server supplies the data from cache memory instead of from disk. The Least Recently Used (LRU) algorithm determines whether data is kept in cache memory or written to disk.

File cache buffers are organized into a linked list. Most Recently Used (MRU) buffers are at the head of the list and least recently used (LRU) buffers are at the end of the list. The length of time the oldest buffer has been in the list is shown as the LRU Sitting Time statistic.

In tuning file cache for the traditional file system, the goal is to determine how much memory the server needs so that recently accessed data can always be retrieved from cache. Monitoring the LRU Sitting Time statistic is key to tuning file cache.

### Understanding LRU Sitting Time

The LRU Sitting Time status is updated and displayed once per second. The statistic is calculated by taking the difference between the current time and the time stamp of the last block in the cache list. (The time stamp is the time when the last block entered the list.) The result is displayed in dd:hh:mm:ss.ts format. The LRU Sitting Time measures the length of time it takes for a cache buffer at the head of the list to make its way down to the tail, where it becomes the least recently used (LRU) buffer. Every cache buffer in the list is being reused within that period of time.

In configurations with excessive cache, the LRU Sitting Time can be very high, even many hours. In configurations with too little cache, the LRU Sitting Time can be measured in seconds. The time can vary widely, depending on circumstances.

On inactive servers, such as those sitting unused overnight or those in lab environments with long periods of idle time, the LRU Sitting Time statistic is incremented by one second every second. This is because the LRU Sitting Time indicates the age of the LRU cache buffer. The LRU Sitting Time statistic is useless under these circumstances, except to confirm the obvious: that data is not being read from or written to the server's cache. This statistic is most useful during peak workloads when the server's cache is servicing the greatest number of users and the broadest possible data set.

You can view the LRU Sitting Time statistic using Novell Remote Manager.

In Novell Remote Manager:

- 1 Click View Memory Config under the Manage Server heading in the navigation frame.
- 2 Click File System Cache.
- 3 Look in the Traditional File System Cache Statistics table.

### Calculating Cache Memory Needs

To determine how much file cache your server needs:

- 1 Track server resource utilization statistics.  
Chart the results for daily, weekly, monthly, period-end, and year-end cycles. Identify recurring periods of peak workloads.
- 2 Observe the LRU Sitting Time statistic during peak workload periods of the day.  
Keep a record of the *lowest* LRU Sitting Time during the peak periods for at least one week, or longer if necessary, to see a consistent pattern.
- 3 Based on the observations you made of the server's resource utilization and LRU Sitting Time, determine an average low LRU Sitting Time.
- 4 Tune the cache.

Size the cache in such a way that the server can sustain an average low LRU Sitting Time of at least 12 minutes during peak workloads.

For example, if your server's original average low LRU was 7 minutes, you will need to add enough memory to increase the LRU Sitting Time to an average low of 12 minutes during peak workloads. This added memory increases the likelihood that repeatedly used data will still be cached when the next user request is received.

On the other hand, if the server's original lowest LRU sitting time was 18 minutes, the server has more than adequate cache. You can leave the excess memory in the server for future growth or remove it and use it somewhere else.

### Setting a File Cache Buffer Alert

You can control when the server alerts you to serious decreases in the number of available file cache buffers. In Novell Remote Manager, click Set Parameters under the Manage Server heading in the navigation frame, and then click Traditional File System.

- ♦ **Minimum File Cache Buffers:** Establishes the minimum number of buffers that will always be reserved for file cache and not made available for other processes. When this minimum is reached, the server displays the message `Cache memory allocator exceeded minimum`.
- ♦ **Minimum File Cache Report Threshold:** Instructs the server to send an alert *before* the minimum number of file cache buffers is reached. Set this parameter to the number of buffers above the minimum at which you want the server to send an alert. For example, if the Minimum File Cache Buffers value is 20 and the Minimum File Cache Report Threshold value is 30, the server sends an alert when 50 file cache buffers are left.

## eDirectory

If you suspect that your server is slow in accessing eDirectory™, see “[Improving eDirectory Performance](#)” in the *Novell eDirectory 8.7.3 Administration Guide*.



# Tuning Logical Address Spaces

Tuning logical address spaces can reduce memory fragmentation problems that might occur because there are not enough physical address spaces to map to logical address spaces in the system. When fragmentation occurs, applications might not be able to use existing memory in the system even though plenty is available. Adjusting the recommended parameters can relieve this problem.

You should tune logical address spaces only after the system has been running for a while in its current condition and when the system shows that there are parameters that you should consider changing.

For example, if you just added a new application or patch to your server; or you started using the server in different circumstances (such as an e-mail server, database server, or storage server), or the number of users accessing the server has increased, you should let the server run in this condition for a week or more before you consider changing anything.

After you have run the server in the current conditions for a week or more, you can use the Tune Logical Address page in Novell Remote Manager to check and optimize logical address space usage and reduce memory fragmentation problems.

## Tune Logical Address Space

Recommended Set Parameter Changes			
Set Parameter	Current Value	New Value	
File Cache Maximum Size	1,250,000,000	1,187,000,000	<a href="#">Change</a>

Restart the server after making the changes using the following command:

```
server -u939000000
```

If you want this change to be permanent, insert this command in the appropriate configuration file.

**Caution:**

The recommended changes are for the current server only. After changing parameters, you need to let the server run for a while to adjust to the services being run with the new settings.

To tune logical address spaces:

- 1 Click View Memory Config under the Manage Server heading in the navigation frame.
- 2 Click Tune Logical Address Space.
- 3 Do one of the following:
  - ◆ If you see a page returned similar to the one shown above with recommended SET parameter changes, click the Change link next to the parameter shown and restart the server using the `-u_value` load command recommended.  
  
If you want the change to be permanent, you can also put the load command in the applicable configuration file, such as the `autoexec.ncf` file.
  - ◆ If you see a page that displays the `No Recommend changes at this time` message, do nothing.

## Tuning Directory Cache for the Traditional File System

This information for tuning directory cache applies only to NetWare servers using the traditional NetWare file system. Directory cache is not used by the Novell Storage Services (NSS) file system. You should tune the directory cache only if you are using the traditional file system.

As directory entries are read and operated upon by a user, NetWare caches the entries to make repeated use of an entry more efficient. In a default configuration, NetWare allocates 500 cache buffers of 4 KB each. Each block read into a cache buffer contains 32 entries.

The default number of buffers is appropriate for only a small number of users. The tuning strategies in this section can help you tune the directory cache to meet your network's requirements.

### Three Directory Cache Tuning Strategies

Sizing the directory cache depends largely on the characteristics of the workload the server supports. The key is the frequency and breadth of directory searches, file opens, closes, and creations.

- ◆ A low-use scenario involves any number of users in which a small number of directories are shared or in which each user's activity is limited to a small region of the directory, such as a home directory.
- ◆ A high-use scenario involves any number of users, but user activity spans a very large number of directories and files. An extreme case might be a document-based system in which document searches routinely traverse large portions of a very large directory.

#### Strategy 1: Tuning for Low Use

To handle low memory use, you probably won't need to allocate any more cache than the NetWare operating system does by default. By default, NetWare allocates 500 buffers immediately upon request, followed by a maximum allocation of up to 2000 buffers (about 2 MB). This is sufficient for the majority of low-use scenarios.

#### Strategy 2: Tuning for High Use

To tune for very high use, change the following SET parameters to the values shown using SET commands at the System Console prompt:

```
minimum directory cache buffers = 2000
maximum directory cache buffers = 4000
```

Or using Novell Remote Manager:

- 1 Click Set Parameters under the Manage Server heading in the navigation frame.
- 2 Click Traditional File System, then select the parameter value link that you want to change.
- 3 Type the new value, then click OK.

To specify a maximum number of concurrent writes of directory cache buffers using Novell Remote Manager:

- 1 Click Set Parameters under the Manage Server heading in the navigation frame.
- 2 Click Common File System > Maximum Concurrent Directory Cache Writes *value*.
- 3 Type the new value, then click OK.

The range is from 5 to 500.

### Strategy 3: Tuning for Name Spaces

When a name space is installed on a traditional NetWare volume, the volume's Directory Entry Table (DET) must include an additional directory entry for each file. So, for example, if a volume supports DOS, NFS, and long name spaces, there are three directory entries in the DET for each file instead of just one.

In such a case, a directory cache buffer with 32 directory entries no longer represents 32 files; it represents only 10 files, because each file now requires three directory entries.

This means the efficiency of your directory cache is decreased by a factor equal to the number of name spaces loaded. Therefore, you might need to add more memory if you load multiple name spaces. By carefully tuning the directory cache, you can maintain high performance even under these conditions.

Before adding name spaces, allow your server to operate in its production environment for several weeks. This lets NetWare allocate the appropriate number of directory cache buffers for the native DOS environment.

After this settling-in period, check the Directory Cache Buffer value in the Monitor General Information window. This is the highest number of directory cache buffers allocated since the server was last started.

Use the following formulas to determine the optimum minimum and maximum number of directory cache buffers:

(# of names spaces, including DOS) x (highest number of Directory Cache buffers since server was last started) = minimum cache buffers

(minimum cache buffers) + 100 = maximum cache buffers

For example, if your server's highest number of Directory Cache buffers since the server was last started is 500 and you want to load one additional name space, you would calculate the minimum and maximum values of directory cache buffers as follows:

$$2 \times 500 = 1000$$

$$1000 + 100 = 1100$$

Before loading the additional name space, change the following SET parameter values using Novell Remote Manager or SET commands at the System Console prompt:

minimum directory cache buffers

maximum directory cache buffers

These settings increase the likelihood that repeatedly used directory cache buffers will remain in cache and that those buffers will remain in cache longer, providing the best possible read response times.

### Monitoring the New Settings

After you tune the directory cache, let the server operate for a few weeks to settle in. Then check to see whether the server allocates the additional directory cache buffers that the new settings allow. If the allocated buffers do not increase to near the level allowed, then you know that your users' directory access patterns don't require the additional cache. On the other hand, if your server uses all the cache you made available, your user community's directory access patterns might be more intensive than you anticipated.

# Using Virtual Memory

NetWare provides a virtual memory system that moves data out of memory and into a swap file on disk if the data isn't used frequently. Thus, the virtual memory system ensures that RAM is used more efficiently but with slower overall performance. It lessens the likelihood that low memory conditions will cause a problem with the server. For general information about virtual memory, see [“Virtual Memory” on page 14](#).

One of the most important things you can do to allocate memory is to set the VM Cache Pool Percentage correctly. If you are running on a system with many applications running in protected mode or with many Java applications, set the percentage high. If you are not running in these two conditions, the VM Cache Pool Percentage doesn't matter.

Novell Remote Manager has several pages that can assist you in monitoring and managing the use of the virtual memory system in your server:

To	Do the Following
View the overall performance of the virtual memory system	<ol style="list-style-type: none"><li>1. Click Health Monitor under Diagnose Server in the navigation frame.</li><li>2. Click Virtual Memory Performance in the Server Health table.</li><li>3. View the Virtual Memory Health Information page.</li></ol>
View specific performance of the virtual memory system	<ol style="list-style-type: none"><li>1. Click View Memory Config under the Manage Server heading in the navigation frame.</li><li>2. Click Virtual Memory System on the View Memory Config page.</li><li>3. Click To Virtual Memory Cache Pool to view the statistics for each cache pool type. (For information on each type of statistic listed, see the online help.)</li><li>4. Click To Virtual Memory Statistics to view the following statistics: Current Available Memory, Free Desired Pages, Free Lots Pages, Free Minimum Pages, Memory Available State, Need Free Clean Pages, Need Free Pages, and Upper Page I/O Threshold. (For information on each type of statistic listed, see the online help.)</li></ol>
Change SET parameters associated with the virtual memory system	<ol style="list-style-type: none"><li>1. Click Set Parameters under the Manage Server heading in the navigation frame.</li><li>2. Click Memory in the Category Name table on the Set Parameters page.</li><li>3. Click the value link for each specific virtual memory parameter you want to change.</li><li>4. Type the new value in the Current Value field, then click OK.</li></ol>

## Using Swap Files

You can manage the swapping of data by controlling the swap files where the data is stored on disk.

You can create one swap file per volume. The swap file for volume `sys:` is created by default; you can delete it if necessary or move it to another volume.

Data moved to disk by virtual memory is stored in any available swap file; it does not matter which volume the swap file is on. Generally, you should place swap files on the fastest volume, the one with the most available space, or the one that is used the least.

You can add a swap file to a volume that is not mounted. The file is created when you mount the volume. When you dismount a volume, the swap file is deleted.

To keep a swap file on that volume, you must create the swap file again when the volume is mounted. The exception is the swap file for volume `sys:`, which is created by default.

Swap files are dynamic; they expand and contract as data is swapped in and out of memory.

To create, delete, and manage swap files, access [Novell Remote Manager](#) and complete the tasks in the following table:

To	Do the Following
Learn about the <code>swap</code> command	<ol style="list-style-type: none"><li>1. Click Console Commands under the Manage Server heading in the navigation frame.</li><li>2. Click SWAP Info on the Console Commands page.</li></ol>
Display the Swap File Information page	<ol style="list-style-type: none"><li>1. Click View Memory Config under the Manage Server heading in the navigation frame.</li><li>2. Click Virtual Memory System &gt; To Swap File Information.</li></ol>
Create a swap file on a designated volume	<ol style="list-style-type: none"><li>1. Click View Memory Config under the Manage Server heading in the navigation frame.</li><li>2. Click Virtual Memory System &gt; To Swap File Information.</li><li>3. Click the value link in the Swapping Enabled column for the volume you want to enable.</li><li>4. (Optional) Type the values you want in the Minimum Reserved, Maximum Used for Swapping, and Minimum Volume Free Space fields. (These parameters specify minimum and maximum size of the swap file and the minimum free space that must be left on the volume.)</li><li>5. On the Enable Swapping page, click Enable.</li><li>6. If the file is not on volume <code>sys:</code>, add the commands to create swap files to the <code>autoexec.ncf</code> file so the files are created each time the server is started. The command syntax is <code>swap add volume_name [parameter=value...]</code> Optional parameters are <code>min</code>, <code>max</code>, and <code>min free</code>. For example: <code>swap add vol1 min=3, max=6, min free=2000</code></li></ol>

To	Do the Following
Delete a swap file from a designated volume	<ol style="list-style-type: none"> <li>1. Click View Memory Config under the Manage Server heading in the navigation frame.</li> <li>2. Click Virtual Memory System &gt; To Swap File Information.</li> <li>3. Click the value link in the Swapping Enabled column for the volume you want to delete the Swap file from.</li> <li>4. On the Disable Swapping page, click Disable.</li> </ol> <p><b>IMPORTANT:</b> If you are using protected address spaces, the Java Virtual Machine, or any other application that uses virtual memory, be sure to keep at least one swap file.</p> <p>If the swap file is being used when it is deleted, then the swapped data is moved to another swap file. If there is no other swap file and there is not enough physical memory to move all the swapped data, an error message is displayed and the file is not deleted.</p> <p>By default, a swap file is created on volume sys:. If you do not want a swap file on volume sys:, place the swap del sys command in the startup.ncf file.</p>
Change the parameter values for a swap file on a designated volume	<ol style="list-style-type: none"> <li>1. Click View Memory Config under the Manage Server heading in the navigation frame.</li> <li>2. Click Virtual Memory System &gt; To Swap File Information.</li> <li>3. Click the link for the value you want to change.</li> <li>4. On the Change Swap Parameters page, specify the new values you want and then click Change.</li> </ol>

## Responding to Disk Thrashing

The real value of virtual memory is in using a sufficient supply of memory more efficiently, thereby improving server performance.

If the overall supply of memory is running low, then swapping will occur more often when virtual memory is in use. If memory is extremely low, the system might spend a significant amount of its time swapping memory in and out of disk and have little time to accomplish useful work. This is called *disk thrashing*.

In extremely low memory conditions, NetWare moves all the memory from a protected address space into the swap file, temporarily stopping the modules within the space. After a period of time, NetWare moves the memory back into the protected address space and shuts down another space, moving its memory to the swap file. Without virtual memory, these extremely low memory conditions would cause processes to fail. With virtual memory, the server keeps running, although very slowly.

If the Average Clean Free Available state or the Average Page I/O state is Suspect or Bad on the Virtual Memory Health Information page in Novell Remote Manager, your server is experiencing disk thrashing.

The best way to prevent disk thrashing is to control how much memory is used by the virtual memory system. Set the VM Cache Pool Percentage parameter to a percentage between 0 to 90.

If you set it to more than the virtual memory system needs, it is harmless because the virtual memory system will not take it. But if you set it to less than the virtual memory system needs, you will create disk thrashing.

To set the VM Cache Pool Percentage in Novell Remote Manager:

- 1** Click Set Parameters under the Manage Server heading in the navigation frame.
- 2** Under Set Parameter Control, click View Hidden Set Parameters to On.
- 3** Under Category Name, click Memory > VM Cache Pool Percentage.
- 4** In the Current Value field, type the new value and then click OK.

If your server is experiencing disk thrashing, you can also use the following two hidden set parameters (Memory category) to manage the balance between the virtual memory cache pool and the file system cache pool:

- ◆ VM Cache Pool Free Desired Pages

Range 0 to 2147483647, Default: 256

This parameter specifies the number of pages that the virtual memory system cache needs to operate properly.

- ◆ VM Cache Pool Free Lots of Pages

Range: 0 to 2147483647, Default: 512

This parameter specifies the number of pages that the virtual memory system cache considers to be excessive. Having an available page count higher than this number might result in the pages being given to another cache pool.

You can specify the point at which the server sends an alert to the console because of excessive memory swapping by using the Average Page In Alert Threshold SET parameter (Memory category).

The value is the average number of pages swapped from disk to memory per second, calculated over the last five seconds. The default is 2000.

To change the values for this parameter using Novell Remote Manager:

- 1** Click Set Parameters under the Manage Server heading in the navigation frame, then click Memory in the Category Name table on the Set Parameters page.
- 2** Click the *value* link for the Average Page in Alert Threshold setting.
- 3** Type the new value in the Current Value field, then click OK.

When disk thrashing occurs, the complete solution is to add more RAM. Virtual memory cannot compensate for an overall lack of server memory, although it can prevent processes from failing and allow a server to continue to function.

## Optimizing Garbage Collection

The NetWare server provides a “**garbage collection**” or memory recovery process that periodically collects allocated memory freed by NLM programs and returns it to cache. For general information on garbage collection, see “**Garbage Collection**” on page 11.

Garbage collection responds to pressure from the virtual memory system. (For an explanation of virtual memory, see “**Virtual Memory**” on page 14.)

If pressure on virtual memory is high, then garbage collection happens immediately.

If pressure on virtual memory is low, then garbage collection happens at the regularly scheduled garbage collection interval. This interval is controlled by the Garbage Collection Interval server parameter.

To view or change this interval in the Novell Remote Manager,

- 1** Click Set Parameters under the Manage Server heading in the navigation frame.
- 2** Click Memory > Garbage Collection Interval.
- 3** View or change the interval setting.

Normally, garbage collection happens in the background. File system cache sizes grow when this occurs.

To initiate garbage collection manually for individual modules using Novell Remote Manager:

- 1** Click List Modules under the Manage Applications heading in the navigation frame.
- 2** In the field, type the name of module that you want to garbage collect memory from and then click Load Module.
- 3** Click the garbage collection link for that module.

A page is returned showing the amount of memory that was freed for use.



# A

## Documentation Updates

The following updates were made to this guide for NetWare® 6.5 Support Pack 3:

Location	Change
Chapter 2, “What’s New,” on page 17	Added information about new and updated features in Support Pack 3.
“Using Server Memory Management Tools” on page 19	Updated procedures, tables, and figures applicable to managing server memory using Novell® Remote Manager.
“Tuning Logical Address Spaces” on page 49	Updated procedures for tuning memory used by logical address spaces using Novell Remote Manager.
Overall document	Changed references from NetWare Remote Manager to Novell Remote Manager and updated links to other manuals.

The following updates were made to this guide for NetWare 6.5 Support Pack 2:

Location	Change
Chapter 2, “What’s New,” on page 17	Added information about new and updated features in Support Pack 2.
Chapter 3, “Managing Server Memory,” on page 19	Added figures throughout the chapter of NetWare Remote Manager page that you can use to manage server memory, as well as and information about each page.
“View Memory Config Page” on page 22	Added information about the Tune Logical Space link and about how to manage logical space.
Overall document	Edited the complete document and made the editorial changes.

