

Hitachi Data Ingestor

6.4.7

File System Protocols (CIFS/NFS) Administrator's Guide

This manual describes precautions and notes on using the CIFS service or NFS service in a Hitachi Data Ingestor (HDI) system from CIFS or NFS clients. Before starting to use the CIFS service or NFS service in an HDI system, read this manual to understand how to set up and operate the services.

© 2017- 2019 Hitachi Vantara Corporation. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara Corporation (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara Corporation at https://support.hitachivantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara Corporation.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, IBM, Lotus, MVS, OS/390, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft Corporate Logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found at <https://www.hitachivantara.com/en-us/company/legal.html>.



Contents

Preface.....	xiii
Intended audience.....	xv
Product version.....	xv
Release notes.....	xv
Organization of HDI manuals.....	xv
Abbreviation conventions.....	xvi
Document conventions.....	xix
Convention for storage capacity values.....	xx
Accessing product documentation.....	xx
Getting help.....	xx
Comments.....	xxi
1 Overview of the CIFS Service.....	1-1
Overview of how to use the CIFS service.....	1-2
2 System Configuration When the CIFS Service Is Used.....	2-1
Products supported by the CIFS service.....	2-2
CIFS clients.....	2-2
Active Directory domain controllers.....	2-4
Network configuration.....	2-5
Configuration in which the CIFS client and the HDI node are connected to the same subnetwork.....	2-6
Configuration in which the CIFS client and the HDI node are connected to different subnetworks.....	2-7
When the CIFS service is used on multiple ports.....	2-7
When the DNS is used.....	2-7
When DHCP is used.....	2-8
3 Using File Services Manager To Run the CIFS Service.....	3-1
Procedure for File Services Manager setup.....	3-2
Setting network information and system information.....	3-2
Editing system files directly.....	3-3
Service configuration definition.....	3-3
Changing the CIFS service configuration definition.....	3-3

Changing the CIFS service configuration definition.....	3-3
Setting an authentication mode.....	3-4
Setting up user mapping.....	3-6
Setting the SMB protocol.....	3-7
Setting up auto-reload.....	3-8
Managing a CIFS file share.....	3-10
Creating a CIFS file share.....	3-10
Editing CIFS file share attributes.....	3-11
Setting up quota information.....	3-12
Using CIFS access logs.....	3-12
What to check before collecting CIFS access logs.....	3-12
Estimating log file size.....	3-13
Information written to the CIFS access log.....	3-15
Backing up the most recent CIFS access logs.....	3-17
4 Managing CIFS Client Users.....	4-1
Procedure for managing users.....	4-2
Local user management.....	4-2
Registering information in an NIS server or an LDAP server (for user authentication)	4-2
Overview of functionality.....	4-3
CSV file formats.....	4-4
Script for registering, deleting, or listing users.....	4-5
Script for CIFS group mapping.....	4-6
Note on users managed by an NIS server or an LDAP server (for user authentication).....	4-8
Notes on local users and group registration.....	4-8
User management in a domain.....	4-8
Setting up an LDAP server for user mapping.....	4-8
Precautions when setting up an LDAP server.....	4-9
Precautions when using OpenLDAP to set up an LDAP server.....	4-9
Precautions when using ADAM to set up an LDAP server.....	4-9
Precautions when using Sun Java System Directory Server to set up an LDAP server	4-10
Examples settings for when using OpenLDAP to set up an LDAP server.....	4-11
Creating a schema file.....	4-11
Setting the index directive.....	4-12
Example settings for when using ADAM to set up an LDAP server.....	4-12
Creating a schema file.....	4-12
Setting index.....	4-15
Example settings for when using Sun Java System Directory Server to set up an LDAP server.....	4-15
Creating a schema file.....	4-15
Setting index.....	4-16
Manually registering a user ID and group ID.....	4-17
How to register IDs with Active Directory.....	4-17
Registering a group ID.....	4-17
Registering a user ID.....	4-18
How to register IDs with an LDAP server.....	4-20
Registering a group ID.....	4-21
Registering a user ID.....	4-21
How to delete IDs registered with an LDAP server.....	4-22

User management when using the RFC 2307 schema.....	4-22
Accessing CIFS shares when an HDI system is accessed from multiple domains.....	4-24
5 User Authentication for CIFS Clients.....	5-1
Local authentication.....	5-2
NT domain authentication.....	5-2
Active Directory authentication.....	5-3
Authentication when user mapping is being used.....	5-7
6 Procedure for Migrating User Resources in a Windows Domain Environment	6-1
.....	6-1
Before performing resource migration.....	6-2
Using the backup utility to perform migration.....	6-7
7 Accessing CIFS Shares.....	7-1
Access method.....	7-2
Notes on access from CIFS client.....	7-3
Notes on CIFS access in an environment where Anti-Virus Enabler is applied.....	7-10
Setting home drives.....	7-11
What is the function for automatically creating a home directory?.....	7-12
Before using the function for automatically creating a home directory.....	7-12
Using home drives.....	7-14
Notes on using the Windows roaming user profile functionality.....	7-15
8 Files and Folders in a CIFS Share.....	8-1
About file and directory names.....	8-2
Supported characters.....	8-2
Notes on the maximum lengths of file names and directory names.....	8-3
Accessing files or directories from a CIFS client.....	8-3
Accessing files or directories from linked functions via a CIFS client.....	8-4
MS-DOS file names in 8.3 format.....	8-4
Notes concerning display of a CIFS share name.....	8-4
Owner or group that owns a file or directory.....	8-5
Access Control Lists.....	8-5
Differences between Classic ACLs and Advanced ACLs.....	8-6
Classic ACL type of file system.....	8-7
Procedure for specifying ACL settings from a CIFS client.....	8-9
How to specify or view the ACL settings for a file.....	8-11
How to specify or view the ACL settings for a folder.....	8-12
Inheriting access permissions from the parent folder.....	8-17
Adding user ACLs or group ACLs.....	8-20
ACL set for a newly created file.....	8-22
ACL set for a newly created folder.....	8-23
SACL.....	8-23
Invalid ACE.....	8-23
Mapping ACL specifications in Windows to file permissions in the HDI system	8-23
.....	8-23
Advanced ACL type of file system.....	8-24
Setting and displaying an ACL from a CIFS client.....	8-24

File system root ACL.....	8-27
ACL-related values.....	8-28
ACL evaluation.....	8-33
ACL initial values, inheritance, and propagation.....	8-34
ACE duplication check.....	8-34
SACL.....	8-35
Invalid ACE.....	8-35
File owners and UNIX permissions.....	8-35
Maximum number of ACL entries that can be set.....	8-37
Migrating to an Advanced ACL type of file system.....	8-37
ACL set by default if there is no inherited ACL.....	8-38
Notes on the case of migrating from Windows.....	8-39
Changing file attributes.....	8-40
Notes on ACLs set by default for new folders and files created in CIFS shares	8-40
Adding user ACLs or group ACLs.....	8-44
File attributes.....	8-47
Setting and checking file attributes from a CIFS client.....	8-47
Whether file attributes can be set.....	8-47
Notes on sharing a file or directory with NFS.....	8-48
Notes on the archive attribute.....	8-49
Notes on the read-only attribute.....	8-49
Offline attribute.....	8-49
Extended attributes in Windows.....	8-50
Timestamps.....	8-51
File access date and time.....	8-51
File modified date and time.....	8-51
File creation date and time.....	8-52
File timestamp resolution.....	8-52
File timestamp management method.....	8-52
File timestamp update resolution.....	8-52
Note for granting file timestamp update permission.....	8-53
Displaying disk capacity.....	8-53
Whether the quotas can be checked on a CIFS client.....	8-55
Disk capacity displayed in accordance with disk usage.....	8-57
Disk capacity displayed when multiple quotas are set.....	8-59
The HDI system.....	8-59
Windows server.....	8-63
WORM files.....	8-64
Access Control by using ABE.....	8-66
How ABE controls whether to display files and folders.....	8-66
About Read permission required for displaying files and folders when ABE is enabled	8-68
Restrictions on files and folders on CIFS shares.....	8-69

9 MMC Linkage.....	9-1
Linking an HDI system with MMC.....	9-2
Operations required to link with MMC (for system administrators).....	9-3
Linking to MMC (for CIFS administrators).....	9-3
Before using the administrative share.....	9-4
CIFS share management from MMC.....	9-4
Viewing a list of CIFS shares.....	9-4

Creating a CIFS share.....	9-5
Changing CIFS share information.....	9-6
Session management from MMC.....	9-7
Viewing a list of sessions.....	9-7
Closing sessions.....	9-8
Managing open files from MMC.....	9-9
List of open files.....	9-9
Closing open files.....	9-9
Share-level ACLs.....	9-10
Notes on using MMC.....	9-12

10 Making Past Versions of Files Available by Using Volume Shadow Copy

Service.....	10-1
Overview of Volume Shadow Copy Service.....	10-2
Compatible CIFS client platforms for Volume Shadow Copy Service.....	10-3
Notes on using Volume Shadow Copy Service on a CIFS client.....	10-3

11 CIFS Client Platforms.....11-1

Notes common to all supported types of Windows.....	11-2
Notes for Windows Server 2008.....	11-2
Files and folders in shared directories.....	11-2
When adding an ACL.....	11-2
When using quotas.....	11-3
When enabling offline files.....	11-3
When using a network drive.....	11-3
When using MMC.....	11-3
Logging on to Windows.....	11-3
Share-level ACLs.....	11-3
When using a large 1-MB MTU.....	11-4
Notes when attempting to access the CIFS service.....	11-4
Notes when you are accessing the CIFS service.....	11-4
Notes for Windows 7.....	11-4
Files and folders in shared directories.....	11-4
When adding an ACL.....	11-5
When using quotas.....	11-5
When using network drives.....	11-5
When enabling offline files.....	11-5
When using MMC.....	11-6
Logging on to Windows.....	11-6
Share-level ACLs.....	11-6
When using a large 1-MB MTU.....	11-6
Notes for Windows 8.....	11-6
Files and folders in shared directories.....	11-7
When adding an ACL.....	11-7
When using quotas.....	11-7
When enabling offline files.....	11-7
When using MMC.....	11-7
Logging on to Windows.....	11-7
Share-level ACLs.....	11-8
Notes for Windows 10 or Windows Server 2016.....	11-8
Files and folders in shared directories.....	11-8

When adding an ACL.....	11-8
When using quotas.....	11-9
When enabling offline files.....	11-9
When using MMC.....	11-9
Logging on to Windows.....	11-9
Share-level ACLs.....	11-9
Notes on accessing the CIFS service.....	11-10
Connecting by using the SMB 1.0 protocol.....	11-10
Notes for Windows Server 2012.....	11-10
Files and folders in shared directories.....	11-10
When adding an ACL.....	11-11
When using quotas.....	11-11
When using MMC.....	11-11
Logging on to Windows.....	11-11
Share-level ACLs.....	11-11
Notes on accessing the CIFS service.....	11-12
Notes for Mac OS X.....	11-12
Support range.....	11-12
Notes on file and directory names.....	11-12
Notes on operations.....	11-13
Notes for Mac OS X v10.9.....	11-14
Notes for Mac OS X v10.10, v10.11, or macOS v10.12.....	11-15

12 Overview of the NFS Service.....	12-1
Overview of using the NFS service.....	12-2

13 System Configuration When the NFS Service Is Used.....	13-1
Products supported by the NFS service.....	13-2
NFS clients.....	13-2
KDC server.....	13-4
ID mapping server.....	13-5
Network configurations.....	13-5
Network configuration when only the NFS service is running.....	13-5
Network configuration when both the CIFS and NFS services are running.....	13-6
Configuring an NFS environment when Kerberos authentication and an NFSv4 domain configuration are used.....	13-7
Configuring an NFS environment when only the NFS service is running.....	13-8
Configuring the KDC server and creating a keytab file.....	13-9
Transferring and installing the keytab file.....	13-9
Set the service configuration definition and create an NFS share from an HDI node.....	13-10
Mounting from an NFS client.....	13-10
Configuring an NFS environment when the CIFS and NFS services are both running at the same time.....	13-10
Creating a keytab file.....	13-11
Transferring and installing the keytab file.....	13-11
Set the service configuration definition and create an NFS share from an HDI node.....	13-12
Mounting from an NFS client.....	13-12

14 Using File Services Manager To Run the NFS Service.....	14-1
File Services Manager setup.....	14-2
Configuring network and system information.....	14-2
Editing system files directly.....	14-3
Service configuration definition.....	14-3
Changing the NFS service configuration definition.....	14-4
Managing NFS shares.....	14-5
Creating an NFS share and changing the settings.....	14-5
Modifying NFS shares.....	14-6
15 Managing NFS Client Users.....	15-1
User management methods.....	15-2
User management when an NFSv4 domain has been set up.....	15-2
16 User Authentication for NFS Clients.....	16-1
User authentication methods.....	16-2
UNIX (AUTH_SYS) authentication.....	16-2
Kerberos authentication.....	16-2
17 Accessing NFS Shares.....	17-1
Access method.....	17-2
Mounting and viewing a file system.....	17-2
When mounting shared directories.....	17-2
When mounting the root directory.....	17-4
Notes on using a file system from an NFS client.....	17-6
Notes on mounting a file system.....	17-6
Notes on using file locking.....	17-7
Notes on using a file system.....	17-10
18 Files and Directories in an NFS Share.....	18-1
File and directory names.....	18-2
ACLs.....	18-2
File attributes.....	18-2
WORM files.....	18-3
Displaying disk capacity.....	18-4
19 Notes on Using File Shares.....	19-1
Notes on accessing file shares.....	19-2
Notes on modifying directories.....	19-3
Managing users who access file shares.....	19-4
Notes on sharing files and directories among CIFS, NFS, and FTP clients.....	19-5
A Troubleshooting when using the CIFS service.....	A-1
syslog.....	A-2
CIFS logs.....	A-3
log.smbd.....	A-3
log.winbindd.....	A-5
MMC operation errors and corrective actions.....	A-8

Errors occurring when a share is added.....	A-9
Errors occurring when the property of a share is changed.....	A-11
Errors occurring when a share is removed.....	A-12
Stopping a share fails due to access denial.....	A-12
Disconnecting a session fails due to access denial.....	A-13
Errors that occur when an open file is closed.....	A-14
Error occurring when a session is displayed.....	A-14
File operation errors and corrective actions.....	A-15
FAQ.....	A-22
My system performance sometimes suffers when CIFS file shares are accessed. Is it possible to improve the system performance?.....	A-22
Is there a user account that is similar to a Windows Administrator account? If so, how can I set one up?.....	A-23
Can only Direct Hosting of SMB be used for the CIFS service?.....	A-23
How can CIFS clients view the Security tab, which allows them to set up or view ACLs?.....	A-23
Can I specify access permissions for entire file systems?.....	A-23
Sometimes it takes time to access the CIFS shares. What is the potential cause of this problem?.....	A-24
The error message "Cannot access file" was displayed when I attempted to access a file in a CIFS share while the on-access scan function of the scan software was enabled.....	A-24
A SID, not the user name or group name, is displayed in the security tab of the properties window of a file or folder. What causes this problem to occur?.....	A-24
Microsoft office files that were correctly overwritten and saved on a CIFS client are displayed as temporary files (.tmp) on other CIFS clients. What causes this problem to occur?.....	A-25
I see the phenomenon that the file I just created gets invisible or the file I just deleted is still visible in the share. What is the cause of these phenomena?.....	A-25

B Troubleshooting when using the NFS service.....B-1

Kerberos authentication errors.....	B-2
Errors in an NFSv4 domain configuration.....	B-4

C Configuring an NFS environment for Kerberos authentication.....C-1

NFS environment to be configured in this appendix.....	C-2
Configuring the KDC server and adding NFS service principals.....	C-3
Before configuring the KDC server.....	C-3
For Windows Server 2008.....	C-3
For Red Hat Enterprise Linux Advanced Platform v5.2.....	C-7
For Solaris 10.....	C-10
For HP-UX 11i v3.....	C-13
For AIX 5L V5.3.....	C-16
Distributing and retrieving keytab files.....	C-18
Keytab file distribution destinations.....	C-18
Distributing keytab files.....	C-18
Retrieving keytab files (for HDI nodes).....	C-19
Retrieving keytab files (for an NFS client).....	C-19

D	Accessing NFS shared directories when Kerberos authentication is used.	D-1
	Specifying a security flavor from File Services Manager.....	D-2
	Mounting shared directories from NFS clients.....	D-2
	Accessing NFS shared directories.....	D-3
E	Adding a secondary KDC server.....	E-1
	Procedure for adding a KDC server.....	E-2
F	APIs for WORM operation.....	F-1
	Creating a WORM file from a CIFS share file.....	F-2
	Creating a WORM file.....	F-2
	APIs required for creating WORM files.....	F-2
	SetFileTime.....	F-2
	SetFileAttributes.....	F-3
	Useful APIs for creating WORM files.....	F-4
	Sample program.....	F-4
	Creating a WORM file from an NFS share file.....	F-6
	Creating a WORM file.....	F-6
	APIs required for creating WORM files.....	F-6
	utime(), utimes().....	F-7
	chmod(), fchmod().....	F-7
	Sample program.....	F-8
	WORM errors and system calls when a WORM file is accessed.....	F-10
G	References.....	G-1
	Web sites.....	G-2
H	Acronyms.....	H-1
	Acronyms used in the HDI manuals.....	H-2

Index



Preface

This manual describes precautions and notes on using the CIFS service or NFS service in a Hitachi Data Ingestor (HDI) system from CIFS or NFS clients.

Before starting to use the CIFS service or NFS service in an HDI system, read this manual to understand how to set up and operate the services.

Please keep this manual in a location that is easily accessible from computers that use the CIFS service or NFS service.

The manual primarily covers the following programs:

- Hitachi File Services Manager
- Configuration Manager

In this manual, all HDI GUI descriptions and operations are for cluster configurations.

This preface includes the following information:

- [Intended audience](#)
- [Product version](#)
- [Release notes](#)
- [Organization of HDI manuals](#)
- [Abbreviation conventions](#)
- [Document conventions](#)
- [Convention for storage capacity values](#)
- [Accessing product documentation](#)

[Getting help](#)

[Comments](#)

Intended audience

This manual is intended for system administrators who are responsible for managing the CIFS or NFS service.

Users are expected to have read all the HDI manuals, such as the *Installation and Configuration Guide* and are expected to also have the following knowledge:

- A basic knowledge of storage systems
- A basic knowledge of Hitachi Content Platform (HCP) systems
- A basic knowledge of networks
- A basic knowledge of file sharing services
- A basic knowledge of SAN
- A basic knowledge of CIFS
- A basic knowledge of NFS
- A basic knowledge of UNIX
- A basic knowledge of Windows
- A basic knowledge of Web browsers

Product version

This document revision applies to Hitachi Data Ingestor version 4.2.1 or later.

Release notes

Read the release notes before installing and using this product. They may contain requirements or restrictions that are not fully described in this document or updates or corrections to this document. Release notes are available on Hitachi Vantara Support Connect: <https://knowledge.hitachivantara.com/Documents>.

Organization of HDI manuals

HDI manuals are organized as shown below.

Note that whether HDI nodes can be set up in a redundant configuration depends on the HDI model. A configuration where nodes are made redundant is called a cluster configuration, and a configuration where a node is not made redundant with another node is called a single-node configuration. Which manuals you need to read depends on which configuration you are going to use.

Manual name	Description
<i>Hitachi Data Ingestor Installation and Configuration Guide, MK-90HDICOM002</i>	You must read this manual first to use an HDI system. This manual contains the information that you must be aware of before starting HDI system operation, as well as the environment settings for an external server.
<i>Hitachi Data Ingestor Cluster Getting Started Guide, MK-90HDICOM001</i>	This manual explains how to set up an HDI system in a cluster configuration.
<i>Hitachi Data Ingestor Cluster Administrator's Guide, MK-90HDI038</i>	This manual provides procedures for using HDI systems in a cluster configuration, as well as provides GUI references.
<i>Hitachi Data Ingestor Cluster Troubleshooting Guide, MK-90HDI029</i>	This manual provides troubleshooting information for HDI systems in a cluster configuration.
<i>Hitachi Data Ingestor Single Node Getting Started Guide, MK-90HDI028</i>	This manual explains how to set up an HDI system in a single-node configuration.
<i>Hitachi Data Ingestor Single Node Administrator's Guide, MK-90HDI039</i>	This manual explains the procedures for using HDI systems in a single-node configuration, as well as provides GUI references.
<i>Hitachi Data Ingestor Single Node Troubleshooting Guide, MK-90HDI030</i>	This manual provides troubleshooting information for HDI systems in a single-node configuration.
<i>Hitachi Data Ingestor CLI Administrator's Guide, MK-90HDI034</i>	This manual describes the syntax of the commands that can be used for HDI systems in a cluster configuration or a single-node configuration.
<i>Hitachi Data Ingestor API References, MK-90HDI026</i>	This manual explains how to use the API for HDI systems in a cluster configuration or a single-node configuration.
<i>Hitachi Data Ingestor Error Codes, MK-90HDI005</i>	This manual contains messages for HDI systems in a cluster configuration or a single-node configuration.
<i>Hitachi Data Ingestor File System Protocols (CIFS/NFS) Administrator's Guide (This manual)</i>	This manual contains the things to keep in mind before using the CIFS or NFS service of an HDI system in a cluster configuration or a single-node configuration from a CIFS or NFS client.

Abbreviation conventions

This manual uses the following abbreviations for product names:

Abbreviation	Full name or meaning
Active Directory	Active Directory(R)
ADAM	Active Directory(R) Application Mode 1.0
File Services Manager	A generic name for the following: <ul style="list-style-type: none"> • Configuration Manager

Abbreviation	Full name or meaning
	<ul style="list-style-type: none"> • Hitachi File Services Manager
HCP	Hitachi Content Platform
HDI	Hitachi Data Ingestor
Mac OS X	<p>A generic name for the following:</p> <ul style="list-style-type: none"> • Mac OS(R) X v10.5 • Mac OS(R) X v10.7 • OS X v10.8 • OS X v10.9 • OS X v10.10 • OS X v10.11 • macOS v10.12
Mac OS X v10.8	OS X v10.8
Mac OS X v10.9	OS X v10.9
Mac OS X v10.10	OS X v10.10
Mac OS X v10.11	OS X v10.11
OpenLDAP	OpenLDAP 2.x
Solaris	<p>A generic name for the following:</p> <ul style="list-style-type: none"> • Solaris 9 Operating System for SPARC Platforms • Solaris 10 Operating System for SPARC Platforms
Solaris 10	Solaris 10 Operating System for SPARC Platforms
Sun Java System Directory Server	Sun Java(TM) System Directory Server 5.2
VMware ESX	VMware vSphere(R) ESX
VMware ESXi	VMware vSphere(R) ESXi(TM)
Windows	Microsoft(R) Windows(R) Operating System
Windows 7	<p>A generic name for the following:</p> <ul style="list-style-type: none"> • Microsoft(R) Windows(R) 7 Enterprise x64 Edition • Microsoft(R) Windows(R) 7 Professional • Microsoft(R) Windows(R) 7 Professional x64 Edition • Microsoft(R) Windows(R) 7 Ultimate
Windows 8	<p>A generic name for the following:</p> <ul style="list-style-type: none"> • Microsoft(R) Windows(R) 8 32-bit • Microsoft(R) Windows(R) 8 64-bit • Microsoft(R) Windows(R) 8 Enterprise 32-bit • Microsoft(R) Windows(R) 8 Enterprise 64-bit • Microsoft(R) Windows(R) 8 Pro 32-bit • Microsoft(R) Windows(R) 8 Pro 64-bit

Abbreviation	Full name or meaning
	<ul style="list-style-type: none"> • Microsoft(R) Windows(R) 8.1 32-bit • Microsoft(R) Windows(R) 8.1 64-bit • Microsoft(R) Windows(R) 8.1 Enterprise 32-bit • Microsoft(R) Windows(R) 8.1 Enterprise 64-bit • Microsoft(R) Windows(R) 8.1 Pro 32-bit • Microsoft(R) Windows(R) 8.1 Pro 64-bit
Windows 10	<p>A generic name for the following:</p> <ul style="list-style-type: none"> • Microsoft(R) Windows(R) 10 Education 32-bit • Microsoft(R) Windows(R) 10 Education 64-bit • Microsoft(R) Windows(R) 10 Enterprise 32-bit • Microsoft(R) Windows(R) 10 Enterprise 64-bit • Microsoft(R) Windows(R) 10 Home 32-bit • Microsoft(R) Windows(R) 10 Home 64-bit • Microsoft(R) Windows(R) 10 Pro 32-bit • Microsoft(R) Windows(R) 10 Pro 64-bit
Windows Server 2008	<p>A generic name for the following:</p> <ul style="list-style-type: none"> • Microsoft(R) Windows Server(R) 2008 Enterprise • Microsoft(R) Windows Server(R) 2008 Enterprise 32-bit • Microsoft(R) Windows Server(R) 2008 Standard • Microsoft(R) Windows Server(R) 2008 Standard 32-bit • Microsoft(R) Windows Server(R) 2008 R2 Enterprise • Microsoft(R) Windows Server(R) 2008 R2 Standard
Windows Server 2008 R2	<p>A generic name for the following:</p> <ul style="list-style-type: none"> • Microsoft(R) Windows Server(R) 2008 R2 Enterprise • Microsoft(R) Windows Server(R) 2008 R2 Standard
Windows Server 2012	<p>A generic name for the following:</p> <ul style="list-style-type: none"> • Microsoft(R) Windows Server(R) 2012 Datacenter • Microsoft(R) Windows Server(R) 2012 Essentials • Microsoft(R) Windows Server(R) 2012 Foundation • Microsoft(R) Windows Server(R) 2012 Standard • Microsoft(R) Windows Server(R) 2012 R2 Datacenter • Microsoft(R) Windows Server(R) 2012 R2 Essentials • Microsoft(R) Windows Server(R) 2012 R2 Foundation • Microsoft(R) Windows Server(R) 2012 R2 Standard
Windows Server 2012 R2	<p>A generic name for the following:</p> <ul style="list-style-type: none"> • Microsoft(R) Windows Server(R) 2012 R2 Datacenter • Microsoft(R) Windows Server(R) 2012 R2 Essentials • Microsoft(R) Windows Server(R) 2012 R2 Foundation • Microsoft(R) Windows Server(R) 2012 R2 Standard
Windows Server 2016	<p>A generic name for the following:</p> <ul style="list-style-type: none"> • Microsoft(R) Windows Server(R) 2016 Datacenter

Abbreviation	Full name or meaning
	<ul style="list-style-type: none"> Microsoft(R) Windows Server(R) 2016 Standard

This manual assumes that the user interface used for Windows operations is from Windows 7 and earlier versions, unless otherwise indicated. If you are using a new user interface from Windows Server 2012 or a later version, refer to the documentation on how to use the new user interface and then, as you read the manual, replace any descriptions of the user interface with those appropriate to your new user interface.

If you want to reference other manuals, note that hereinafter in this manual, the *Hitachi Data Ingestor Cluster Administrator's Guide* and *Hitachi Data Ingestor Single Node Administrator's Guide* are referred to as the *Administrator's Guide*, and the *Hitachi Data Ingestor Cluster Troubleshooting Guide* and the *Hitachi Data Ingestor Single Node Troubleshooting Guide* are referred to as the *Troubleshooting Guide*. See the appropriate manual as needed.

Document conventions

This document uses the following typographic conventions:

Convention	Description
Bold	Indicates text on a window, other than the window title, including menus, menu options, buttons, fields, and labels. Example: Click OK .
<i>Italic</i>	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: <i>copy source-file target-file</i> <i>Note:</i> Angled brackets (< >) are also used to indicate variables.
screen/code	Indicates text that is displayed on screen or entered by the user. Example: # pairdisplay -g oradb
< > angled brackets	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: # pairdisplay -g <group> <i>Note:</i> Italic font is also used to indicate variables.
[] square brackets	Indicates optional values. Example: [a b] indicates that you can choose a, b, or nothing.
{ } braces	Indicates required or expected values. Example: { a b } indicates that you must choose either a or b.
vertical bar	Indicates that you have a choice between two or more options or arguments. Examples: [a b] indicates that you can choose a, b, or nothing. { a b } indicates that you must choose either a or b.
<u>underline</u>	Indicates the default value. Example: [<u>a</u> b]

Convention	Description
...	The item or items preceding the ellipsis (...) can be repeated. To specify multiple items, use a comma (,) to delimit them. Example: A,B... indicates that B can be specified as many times as necessary after A.

Convention for storage capacity values

Storage capacity values (e.g., drive capacity) are calculated based on the following values:

Capacity Unit	Physical Value	Logical Value
1 KB	1,000 bytes	1,024 (2^{10}) bytes
1 MB	1,000 KB or $1,000^2$ bytes	1,024 KB or $1,024^2$ bytes
1 GB	1,000 MB or $1,000^3$ bytes	1,024 MB or $1,024^3$ bytes
1 TB	1,000 GB or $1,000^4$ bytes	1,024 GB or $1,024^4$ bytes
1 PB	1,000 TB or $1,000^5$ bytes	1,024 TB or $1,024^5$ bytes
1 EB	1,000 PB or $1,000^6$ bytes	1,024 PB or $1,024^6$ bytes
1 block	-	512 bytes

Accessing product documentation

Product user documentation is available on Hitachi Vantara Support Connect: <https://knowledge.hitachivantara.com/Documents>. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

Getting help

[Hitachi Vantara Support Connect](#) is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to Hitachi Vantara Support Connect for contact information: https://support.hitachivantara.com/en_us/contact-us.html.

[Hitachi Vantara Community](#) is a global online community for Hitachi Vantara customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. **Join the conversation today!** Go to community.hitachivantara.com, register, and complete your profile.

Comments

Please send us your comments on this document to doc.comments@hitachivantara.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara Corporation.

Thank you!

Overview of the CIFS Service

CIFS clients can access data via the CIFS service in a Hitachi Data Ingestor (HDI) system. This chapter provides an overview of using the CIFS service.

- [Overview of how to use the CIFS service](#)

Overview of how to use the CIFS service

If the system administrator creates CIFS shares in a file system or in a directory, CIFS clients can access data in the storage system over the network.

The following figure shows the flow of how a CIFS client accesses data in the file system:




Figure 1-1 Flow of how a CIFS client accesses data in the file system

System Configuration When the CIFS Service Is Used

This chapter describes what kinds of HDI system configurations are required for using the CIFS service and also the environments in which the CIFS service can run.

- [Products supported by the CIFS service](#)
- [Network configuration](#)

Products supported by the CIFS service

The CIFS service supports the following products.

CIFS clients

The following client OSs can be used for the CIFS service:

Table 2-1 Products that support CIFS clients

CIFS client platform	CIFS connectivity					
	IPv4 connections	IPv6 connections	SMB 1.0	SMB 2.0	SMB 2.1	SMB 3.0
Mac OS(R) X v10.5	Y	N	Y	N	N	N
Mac OS(R) X v10.7	Y	Y	Y	N	N	N
Mac OS(R) X v10.8	Y	Y	Y	N	N	N
Mac OS(R) X v10.9	Y	Y	Y	N	N	N
Mac OS(R) X v10.10	Y	Y	N	Y	N	N
Mac OS(R) X v10.11	Y	Y	N	Y	N	N
macOS v10.12	Y	Y	N	Y	N	N
Microsoft(R) Windows(R) 7 Enterprise x64 Edition (with SP1)	Y	Y	Y	Y	Y	N
Microsoft(R) Windows(R) 7 Professional (with SP1)	Y	Y	Y	Y	Y	N
Microsoft(R) Windows(R) 7 Professional x64 Edition (with SP1)	Y	Y	Y	Y	Y	N
Microsoft(R) Windows(R) 7 Ultimate (with SP1)	Y	Y	Y	Y	Y	N
Microsoft(R) Windows(R) 8 32-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 8 64-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 8 Enterprise 32-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 8 Enterprise 64-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 8 Pro 32-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 8 Pro 64-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 8.1 32-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 8.1 64-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 8.1 Enterprise 32-bit	Y	Y	Y	Y	Y	Y

CIFS client platform	CIFS connectivity					
	IPv4 connections	IPv6 connections	SMB 1.0	SMB 2.0	SMB 2.1	SMB 3.0
Microsoft(R) Windows(R) 8.1 Enterprise 64-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 8.1 Pro 32-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 8.1 Pro 64-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 10 Education 32-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 10 Education 64-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 10 Enterprise 32-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 10 Enterprise 64-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 10 Home 32-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 10 Home 64-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 10 Pro 32-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows(R) 10 Pro 64-bit	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows Server(R) 2008 Enterprise (with SP2)	Y	Y	Y	Y	N	N
Microsoft(R) Windows Server(R) 2008 Enterprise 32-bit (with SP2)	Y	Y	Y	Y	N	N
Microsoft(R) Windows Server(R) 2008 Standard (with SP2)	Y	Y	Y	Y	N	N
Microsoft(R) Windows Server(R) 2008 Standard 32-bit (with SP2)	Y	Y	Y	Y	N	N
Microsoft(R) Windows Server(R) 2008 R2 Enterprise (with SP1)	Y	Y	Y	Y	Y	N
Microsoft(R) Windows Server(R) 2008 R2 Standard (with SP1)	Y	Y	Y	Y	Y	N
Microsoft(R) Windows Server(R) 2012 Datacenter	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows Server(R) 2012 Essentials	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows Server(R) 2012 Foundation	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows Server(R) 2012 Standard	Y	Y	Y	Y	Y	Y

CIFS client platform	CIFS connectivity					
	IPv4 connections	IPv6 connections	SMB 1.0	SMB 2.0	SMB 2.1	SMB 3.0
Microsoft(R) Windows Server(R) 2012 R2 Datacenter	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows Server(R) 2012 R2 Essentials	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows Server(R) 2012 R2 Foundation	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows Server(R) 2012 R2 Standard	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows Server(R) 2016 Datacenter	Y	Y	Y	Y	Y	Y
Microsoft(R) Windows Server(R) 2016 Standard	Y	Y	Y	Y	Y	Y

Legend: Y = Connection possible; N = Connection not possible

Note:

For a CIFS client running Windows 8 or later or Windows Server 2012 or later, if the HDI system is configured so that SMB 2.0 is used when CIFS clients access the system, the CIFS client might not be able to access the HDI system. For information on how to fix this problem, see [File operation errors and corrective actions on page A-15](#).

Active Directory domain controllers

The following OSs can be used on the Active Directory domain controllers:

Table 2-2 Products that support Active Directory domain controllers

Active Directory domain controller platform	CIFS connectivity	
	IPv4 connections	IPv6 connections
Microsoft(R) Windows Server(R) 2008 Enterprise (with SP2)	Y	Y
Microsoft(R) Windows Server(R) 2008 Enterprise 32-bit (with SP2)	Y	Y
Microsoft(R) Windows Server(R) 2008 Standard (with SP2)	Y	Y
Microsoft(R) Windows Server(R) 2008 Standard 32-bit (with SP2)	Y	Y
Microsoft(R) Windows Server(R) 2008 R2 Enterprise (with SP1)	Y	Y

Active Directory domain controller platform	CIFS connectivity	
	IPv4 connections	IPv6 connections
Microsoft(R) Windows Server(R) 2008 R2 Standard (with SP1)	Y	Y
Microsoft(R) Windows Server(R) 2012 Datacenter	Y	Y
Microsoft(R) Windows Server(R) 2012 Essentials	Y	Y
Microsoft(R) Windows Server(R) 2012 Foundation	Y	Y
Microsoft(R) Windows Server(R) 2012 Standard	Y	Y
Microsoft(R) Windows Server(R) 2012 R2 Datacenter	Y	Y
Microsoft(R) Windows Server(R) 2012 R2 Essentials	Y	Y
Microsoft(R) Windows Server(R) 2012 R2 Foundation	Y	Y
Microsoft(R) Windows Server(R) 2012 R2 Standard	Y	Y
Microsoft(R) Windows Server(R) 2016 Datacenter	Y	Y
Microsoft(R) Windows Server(R) 2016 Standard	Y	Y

Legend: Y = Connection possible; N = Connection not possible

The versions of the SMB protocol that are supported vary according to the platform of the Active Directory domain controller. For this reason, you need to make sure that the Active Directory domain controller you are using supports the versions of the SMB protocol that the CIFS service uses to communicate with the domain controller. If it does not support those versions of the SMB protocol, change the values specified for the options `client_ipc_max_protocol` and `client_ipc_min_protocol`, of the command `cifsoptset`, that determine the versions of the SMB protocol that the CIFS service uses to communicate with the domain controller.

Network configuration

Hitachi Data Ingestor (HDI) supports the CIFS protocols shown in the following table. The CIFS client can use the CIFS service by using the virtual IP address (IPv4 or IPv6), the host name, or the NetBIOS name for HDI nodes.

Table 2-3 CIFS protocols supported in the HDI system

#	Protocol
1	NetBIOS over TCP/IP
2	Direct Hosting of SMB

When HDI is installed as a new installation, in the initial status, the NetBIOS over TCP/IP protocol is disabled to reduce data communication load and

security risk. Change the CIFS service configuration so that the NetBIOS over TCP/IP protocol is used when you want to do any of the following:

- Use the browsing function
Only clients that are connected via IPv4 can use the browsing function.
- Use WINS, `lmhosts`, or broadcasting to resolve an HDI node NetBIOS name from the CIFS client

When you use the browsing function on the network:

- Make sure that nodes in a cluster join the same workgroup, the same NT domain, or the same Active Directory domain.
- When you specify the NetBIOS name of an HDI node from the CIFS client, specify the host name of the node to be accessed.
- The behavior of the HDI system might differ from that described in this manual, depending on the settings and state of the network to which the HDI system is connected. In this case, fix any duplicated network addresses, as well as server settings or router settings, to make sure that the overall network is operating normally.
- The browsing function can only be used by clients that are connected via IPv4.

Configuration notes for various network examples are provided below. These examples assume that when you use the NetBIOS name of a node for the CIFS service in an HDI system, all machines in the network can resolve names through WINS, DNS, `lmhosts`, or other such services. The three network configurations explained are listed below. For a more detailed description of these network configurations, see the *Installation and Configuration Guide*.

- Configuration in which the CIFS client and the HDI node are connected to the same subnetwork
- Configuration in which the CIFS client and the HDI node are connected to different subnetworks
- The CIFS service is used on multiple ports

Configuration in which the CIFS client and the HDI node are connected to the same subnetwork

Note the following when you use the browsing function if both the CIFS client and the HDI node are connected to the same subnetwork:

- On the CIFS client, we recommend that you use the WINS server to perform name resolution.
- When the domain controller exists on a different network, the CIFS service in the HDI system might run as the local master browser. If a failover occurs in this configuration, the CIFS service that has been running as the local master browser will temporarily stop, causing the CIFS client to take longer to obtain a computer list. Make sure that the CIFS client accesses a CIFS share when the CIFS service starts to run as a local master browser.

Configuration in which the CIFS client and the HDI node are connected to different subnetworks

Note the following when you use the browsing function if the CIFS client is connected to a different subnet from the one the HDI node is connected to:

- Be sure to use an NT domain configuration or Active Directory configuration.
- The subnetwork to which an HDI node is connected must contain a domain controller.
- When the WINS server is used as a name server for the CIFS client, we recommend that all CIFS client in the network be set in the WINS client.
- If no WINS server is used, the `lmhosts` file needs to be edited as follows:
 - For NT domain configurations:
Add the following to the `lmhosts` file of the backup domain controller, as well as to all the `lmhosts` files of CIFS clients, for subnets to which no domain controller is connected.
primary-domain-controller-ip-address domain-name#1B
 - For Active Directory domain configurations:
Add the following to the `lmhosts` file of the domain controller on the same subnet as the CIFS client. For subnets not connected to a domain controller, add the following to the `lmhosts` file for all CIFS clients.
IP-address-of-domain-controller-in-subnetwork-that-contains-HDI-node domain-name#1B

When the CIFS service is used on multiple ports

When you use the CIFS service on multiple ports, every subnetwork connected to a port requires a separate WINS server. Any CIFS client connected to a network can select the access path to either the HDI node depending on the WINS server being used.

When the DNS is used

While multiple IP addresses are assigned to a server name of domain controller specified when CIFS service authentication mode is set, or server name of domain controller of a trusted domain, if an IP address that is not accessible from HDI is contained, access to the domain controller may be disabled during operation. This problem is thought to occur because the DNS round-robin function distributes the processing for all of the assigned IP addresses, and as a result, an IP address that cannot be accessed from the HDI system is drawn as the address for the domain controller. Note that you can use the `nslookup` command for Windows to check whether the round-robin function is being used. If a problem like this occurs, you can fix it by registering the domain controller in the `/etc/hosts` file. In the **Edit System File** page of the **Network & System Configuration** dialog, add the following information to the `/etc/hosts` file:

IP-address-that-can-be-accessed-from-the-HDI-system Domain-controller-host-name Domain-controller-host-name (FQDN)

The following is an example of the information to be added:

```
10.213.89.113 dc-host dc-host.sample.domain.local
```

When DHCP is used

When DHCP is used in a single-node configuration, if the DNS resolver cache is enabled on the Windows client, the result of name resolution might be outdated and a CIFS access error might occur. In this case, execute the `ipconfig /flushdns` command on the Windows client to flush the data in the DNS resolver cache, and then try accessing CIFS again.

Using File Services Manager To Run the CIFS Service

The system administrator performs various management tasks in order to properly maintain the HDI system. Within those various management tasks, this chapter describes the operations that require particular attention for the use of the CIFS service. The information in this chapter assumes that the File Services Manager GUI is used.

- [Procedure for File Services Manager setup](#)
- [Setting network information and system information](#)
- [Service configuration definition](#)
- [Managing a CIFS file share](#)
- [Setting up quota information](#)
- [Using CIFS access logs](#)

Procedure for File Services Manager setup

The system administrator uses File Services Manager to configure the information required to start HDI system operations. The following figure shows the setup procedure for File Services Manager. Among the operations shown in the figure below, this manual mainly describes the operations that require particular attention to the use of the CIFS service. For details on the other operations, see the *Administrator's Guide*.

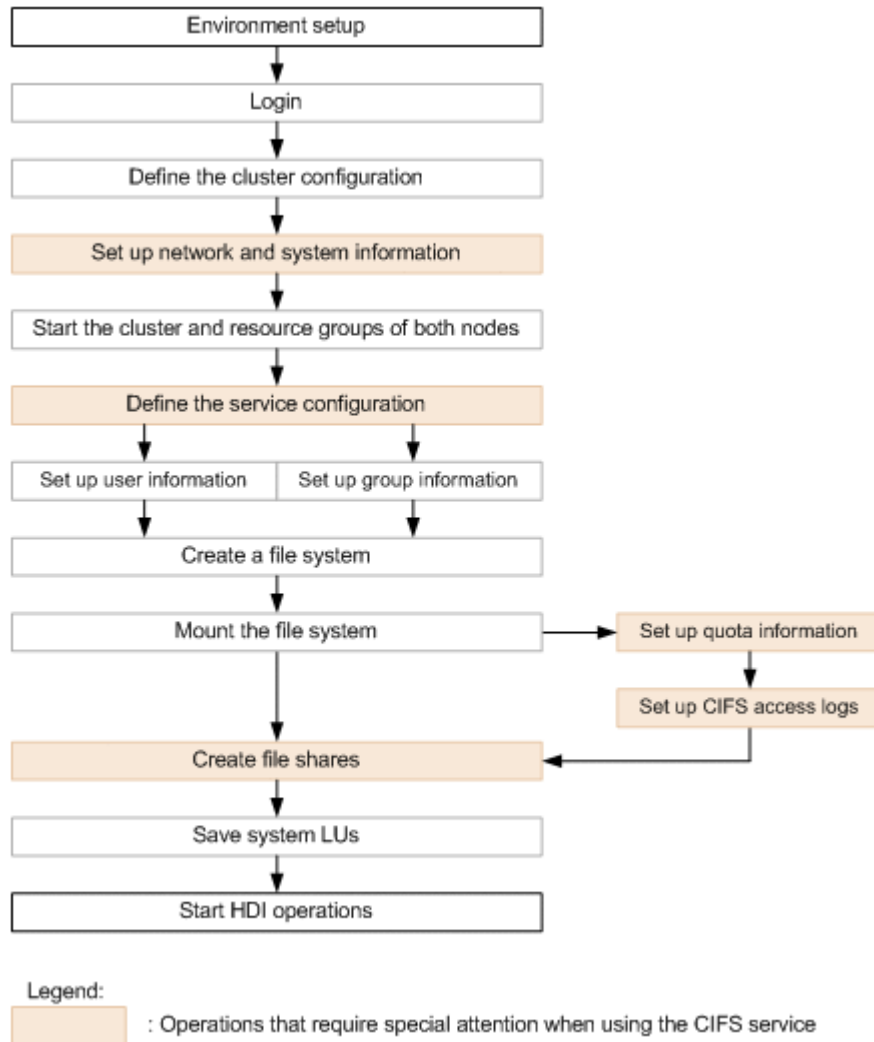


Figure 3-1 File Services Manager configuration procedure

Setting network information and system information

When necessary, the system administrator can specify and change the interface information, network information, and linked external server information for each HDI node in the **System Setup Menu** page of the **Network & System Configuration** dialog box. The following explains precautions to keep in mind when directly editing system files. For details on

the other settings available from the **System Setup Menu** page, see the *Administrator's Guide*.

Editing system files directly

The system administrator can directly edit the system files of the HDI system in the **Edit System File** page of the **Network & System Configuration** dialog box. For details about how to directly edit the system files and the settings that can be specified, see the *Administrator's Guide*. These files must be configured for each node so that the settings of each node match the settings of the other nodes in a cluster.

The following explains the system files to edit and in what cases to edit those system files for CIFS service use:

- /etc/hosts
Edit this system file by specifying host names to limit the CIFS clients that can access a node or CIFS share.
- /etc/cifs/lmhosts
Edit this system file when it is necessary to search for the domain controllers of domains with which trust relationships have been established and the authentication mode of the CIFS service is Active Directory authentication or NT domain authentication.

Service configuration definition

The following table shows the CIFS service contents that can be managed by a system administrator. For details on how to manage this service, see the *Administrator's Guide*.

Table 3-1 Contents managed for the CIFS service

Service type	Service name	Configuration definition changes	Service maintenance	Start / stop / restart
CIFS Service	CIFS	Y	Y	Y

Legend: Y = Possible

Changing the CIFS service configuration definition

This subsection provides supplementary details on making changes to the CIFS service configuration definitions.

Changing the CIFS service configuration definition

For details about how to change the CIFS service configuration definitions and notes on changing the definitions, see the *Administrator's Guide*. This subsection contains additional notes on changing configuration definitions of

the CIFS service in the **CIFS Service Management** page of the **Access Protocol Configuration** dialog box.

Table 3-2 Notes on the CIFS service setup in the CIFS Service Management page (Setting Type: Security)

#	Item	Explanation
1	Host access restrictions	<p>When you specify a network, you can use either of the following formats:</p> <p>For IPv4</p> <p>When specifying a network address: IP address (For example, 10.203.15.0)</p> <p>When specifying a network range according to a netmask: <i>network-address/netmask</i> (For example, 10.203.15.0/255.255.255.0)</p> <p>For IPv6</p> <p>When specifying an address prefix: IP address (For example, fe80::223:7dff:0:0)</p> <p>When specifying a network range according to a prefix length: <i>address-prefix/prefix-length</i> (For example, fe80::223:7dff:0:0/64)</p>

Setting an authentication mode

Three authentication modes can be selected for the CIFS service. For details about how to configure the authentication mode and notes on configuring the authentication mode, see the *Administrator's Guide*. This subsection provides supplementary notes on setting the authentication mode.

Table 3-3 Notes for setting of authentication modes of CIFS service

#	Authentication mode	Description and precautions
1	Local authentication	<p>Information specified in the Local Authentication page is as follows:</p> <ul style="list-style-type: none"> Specify the name of the workgroup that the node belong to. Specify a workgroup name that is different from the host name of the node. If you specified the same name, a group name would not be displayed properly when ACL is set.
2	NT domain authentication	<p>Information specified in the NT Domain Authentication page is as follows:</p> <ul style="list-style-type: none"> For Domain name, specify a NT domain name.

#	Authentication mode	Description and precautions
		<ul style="list-style-type: none"> For Domain administrator name, specify a character string that does not include the marks % or @.
3	Active Directory authentication	<p>Information specified in the Active Directory Authentication page is as follows:</p> <ul style="list-style-type: none"> Specify a DNS name of an Active Directory domain for Domain name. All lower-case letters are interpreted as upper-case letters. <p>In the domain controller policies, when Require signing is set for the Domain controller: LDAP server signing requirements policy, if you do not set the LDAP communication to be signed, attempts to join the domain will fail. Make sure to use the <code>cifsoptset</code> command in advance to set the LDAP communication to be signed.</p> <ul style="list-style-type: none"> You cannot use a percent sign (%) or at mark (@) in Domain user name. <p>For the user name, specify a domain user that does not include any of the characters noted above.</p> <p>The user account specified for Domain user name is used as the user account when a node is added to or removed from the Active Directory domain.</p> <p>If the specified user is a general domain user who does not have administrator privileges for the domain, the following restrictions apply.</p> <ul style="list-style-type: none"> They cannot add more than 10 servers (here, <i>servers</i> refers to HDI nodes) to an Active Directory domain. However, they can add more than 10 servers if they belong to the <code>Account Operators</code> group in the domain. For this reason, if the specified user has already added 10 servers to a domain and needs to add more servers to the domain, verify that the user belongs to the <code>Account Operators</code> group in the domain. When a node is removed from the Active Directory domain, the computer account of the node cannot be deleted from the Active Directory domain. To delete the computer account when a node is removed, make sure that the specified user belongs to one of the following groups, or grant deletion privileges to the specified user: <ul style="list-style-type: none"> - <code>Domain Admins</code> - <code>Account Operators</code> - <code>Enterprise Admins</code> <p>To change a user who has already been set as a domain user to a different user, perform either of the following operations in advance:</p> <ul style="list-style-type: none"> From the domain controller, delete the computer account corresponding to any nodes that were added to the domain by the old user. From the domain controller, add the ACLs of the new user to the computer account corresponding to the

#	Authentication mode	Description and precautions
		<p>target node, and then grant permissions for the following operations to that user:</p> <ul style="list-style-type: none"> - Read - Validated write to DNS host name - Validated write to service principal name - Reset password - Change password - Write account restrictions <p>Note: If setting changes such as the following are made, the user account is removed from the Active Directory.</p> <ul style="list-style-type: none"> o The domain of the Active Directory to which the user account was added is changed. o The user authentication method is changed to a method other than Active Directory authentication. <ul style="list-style-type: none"> • Even though you specify an incorrect value for Domain name (NetBIOS), if the Domain name value is correct, you can successfully start or restart the CIFS service. In this case, if you perform an ACL operation for a shared file, a problem such as not being able to communicate with the domain controller might occur. Specify values carefully to avoid such problems. • Depending on the information that is set for the computer account of the domain that the HDI system is to join, an attempt to join the domain might fail. If a computer account that has the same host name as the HDI system is registered for the domain that HDI system is to join, perform either of the following steps before joining the domain. <ul style="list-style-type: none"> o Change the host name of the HDI system so that the information about the computer account that has the same host name as the HDI system will not be overwritten. o If the computer account that has the same host name as the HDI system does not exist, manually delete that computer account from the domain. <p>In addition, for notes about Active Directory authentication, see Active Directory authentication on page 5-3.</p>

Setting up user mapping

For details on the settings and notes related to user mapping, see the *Administrator's Guide*.

The following provides supplementary notes that apply only when **Use user mapping using LDAP** is selected as the user mapping method:

User mapping functionality cannot be used with an OpenLDAP server using transport layer security (TLS), a protocol for sending and receiving encrypted information over the Internet.

Setting the SMB protocol

For CIFS access, clients use a file sharing protocol called *SMB* to access HDI nodes. In an HDI system, you can select the SMB protocol version to use in the CIFS service configuration definition.

The version of the SMB protocol to be used by a client for CIFS access is determined by the negotiation between the HDI system and the client when a connection is established, based on the platform of the CIFS client and the CIFS service configuration definition. The table below shows the versions of SMB protocols used for CIFS access.

Table 3-4 SMB protocol version used for CIFS access

CIFS client platform	CIFS service configuration definition			
	SMB 1.0	SMB 2.0	SMB 2.1	SMB 3.0
Mac OS X v10.5 Mac OS X v10.7 Mac OS X v10.8	SMB 1.0	SMB 1.0	SMB 1.0	SMB 1.0
Mac OS X v10.9 ^{#1}	SMB 1.0	SMB 1.0	SMB 1.0	SMB 1.0
Mac OS X v10.10 ^{#2} Mac OS X v10.11 ^{#2} macOS v10.12 ^{#2}	--	SMB 2.0	--	--
Windows Server 2008 (excluding Windows Server 2008 R2)	SMB 1.0	SMB 2.0	SMB 2.0	SMB 2.0
Windows 7 Windows Server 2008 R2	SMB 1.0	SMB 2.0	SMB 2.1	SMB 2.1
Windows 8 Windows 10 Windows Server 2012	SMB 1.0	SMB 2.0	SMB 2.1	SMB 3.0
Windows Server 2016	SMB 1.0 ^{#3}	SMB 2.0	SMB 2.1	SMB 3.0

Legend:

--: Not supported

#1

When using a Mac OS X v10.9 client, you must change the settings in the HDI system or on the client so that the SMB 1.0 protocol is used to access CIFS resources. For details about how to change these settings, see [Notes on operations on page 11-13](#)

#2

When using a Mac OS X v10.10, v10.11, or macOS v10.12 client, you must change the settings in the HDI system so that the SMB 2.0 protocol is used to access CIFS resources. For details about how to change these

settings, see [Notes for Mac OS X v10.10, v10.11, or macOS v10.12 on page 11-15](#).

#3

To use the SMB 1.0 protocol to access CIFS resources, you must change the settings in the HDI system or on the client in advance. For details about how to change these settings, see [Connecting by using the SMB 1.0 protocol on page 11-10](#).

If the settings for the version of the SMB protocol via a CIFS service configuration definition change, log off from any CIFS clients that were connected to the CIFS service before the change was made, and then log in again.

Note that the HDI system supports the following for SMB 2.0, SMB 2.1, and SMB 3.0.

SMB 2.0

- Requests for multiple commands within the same packet
- Buffer size expansions
- Increases to the number of files and shares that can be handled by the SMB protocol

SMB 2.1

Large 1-MB MTU

For details about how to enable a large 1-MB MTU on the CIFS client, contact Microsoft Support.

SMB 3.0

SMB encryption

Note that access performance will be reduced if you encrypt CIFS client communication by using SMB encryption, compared to not using any encryption.

Note:

If you enable SMB encryption for a CIFS share, you will not be able to use a client cache, regardless of the CIFS service settings and CIFS share settings.

Setting up auto-reload

In a CIFS service, you can specify whether to automatically reload the CIFS share settings when they are changed. For details on how to set up auto-reload, and for related notes, see the *Administrator's Guide*. This section includes supplementary information about the CIFS share settings that are automatically reloaded if auto-reload is enabled.

Table 3-5 CIFS share settings that are automatically reloaded

#	Settings	Remarks
1	Read-only settings	To change the settings, use the GUI.

#	Settings	Remarks
2	Settings for specially permitted users and groups	For details about the GUI, see the <i>Administrator's Guide</i> .
3	Settings for host-based or network-based access restriction	
4	Browsable share settings	
5	Settings allowing guest account access	
6	Settings to assign access permissions only for the owner	
7	Access permissions for new files	
8	Access permissions for new directories	
9	Name of the CIFS share	
10	The comment shown to CIFS clients	
11	Settings to enable the automatic creation of the home directory	
12	Users who are allowed to change file time stamps	
13	Disk synchronization policy settings	
14	Windows client access policy settings	
15	Settings to allow caching by the CIFS client	
16	Settings to allow read-only client caching during access conflicts	
17	Settings to allow access-based enumeration	
18	Settings for using the Volume Shadow Copy Service	
19	SMB encryption settings	
20	Settings for offline attributes	
21	<code>case_sensitive</code> option settings#	To change the settings, use the <code>cifsoptset</code> command. For details about the <code>cifsoptset</code> command, see the " <i>CLI Administrator's Guide</i> ".
22	<code>hide_system_files</code> option settings#	
23	<code>logging_unsupported_access</code> option settings#	
24	<code>check_strict_allocate</code> option settings#	

#

If you execute a command by specifying `-x` option with `default` specified, log out from the CIFS client and then log in again, or restart the CIFS service, after completing the command execution.

Managing a CIFS file share

This section contains notes on how a system administrator can use File Services Manager to create a CIFS file share, and how they can modify the CIFS file share's attributes.

Creating a CIFS file share

The system administrator can create a CIFS share in the **Add Share** dialog box or in the **Create and Share File System** dialog box. For details on how to create a CIFS share, see the *Administrator's Guide*. The following are additional notes on creating a CIFS share:

When using only the CIFS service to share files and directories, perform the following settings:

- In any of the following dialog boxes, select the **Yes** check box for **Record last access time**: Note that even if **Yes** is not selected, depending on the operating specifications of an application such as Microsoft Excel, the access time of a file might change when the file is updated.
 - In the **Create and Share File System** dialog box, the **Advanced** tab
 - In the **Create File System** dialog box, the **Advanced** tab
 - In the **Mount File System** dialog box
- For a file system of the Classic ACL type, if you want to allow users other than the file owner to change file modification times from the CIFS service, select **Write permitted users** for **Users allowed to change file time stamp** in the following locations:
 - In the **Create and Share File System** dialog box, the **Advanced** tab
 - In the **Add Share** dialog box, the **Advanced** tab
 - In the **Edit Share** dialog box, the **Advanced** tab

For a file system of the Advanced ACL type, the *write attribute* privilege in the ACL determines whether the file modification time can be changed.

Note that for a file system of the Advanced ACL type, you can always manipulate ACLs.

Keep the following in mind when using CIFS and NFS to share files and directories:

In the CIFS service, when users other than the file owner are permitted to change a file update date, all users with write permissions for the file can use a CIFS client to change the file update date. Since NFS clients do not allow users other than the file owner are permitted to change a file

update date, take precautions when sharing the same file with CIFS clients and NFS clients.

If Microsoft Word, Excel, or PowerPoint is used to view or update a file in a CIFS share, turn on Enable ACL option when creating a CIFS share.

If the node contains 64 GB of memory, and if automatic reloading is enabled, the maximum number of CIFS shares that can be created for a single cluster is 1,024 shares. If automatic reloading of CIFS share settings is enabled in the CIFS service configuration, when CIFS shares are created or modified, the CIFS share settings are automatically reloaded and applied to the CIFS client environment. If there are a large number of CIFS client connections, reloading places a heavy load on the HDI system and causes the file access response to deteriorate temporarily. To minimize the deterioration in response, run the system by using the recommended values. Note that file access response immediately after reloading might take as long as 20 seconds.

The recommended values for the number of CIFS clients that can be connected and the number of CIFS shares are listed below.

Table 3-6 The recommended values for the number of CIFS clients that can be connected and the number of CIFS shares

CIFS client connections	CIFS shares
1,600 or fewer	1,024 or fewer
1,601 to 2,000	768 or fewer
2,001 to 2,400	512 or fewer
2,401 to 9,600	256 or fewer

Note: The recommended values are calculated based on performance under the following conditions: the node contains 64 GB of memory, there is one CPU (four cores), and read and write operations for a single text file are performed from the CIFS client every five minutes.

Editing CIFS file share attributes

A system administrator can use the **Edit Share** dialog box to edit attributes for CIFS file shares. For details about how to edit CIFS share attributes and notes on changing the attributes, see the *Administrator's Guide*. The following shows notes for when CIFS file shares are edited.

- The information currently set is applied for items for which information has not changed.

In addition to the above note, see also the notes on creating a CIFS share in [Creating a CIFS file share on page 3-10](#).

Setting up quota information

The system administrator can set quotas for file systems or directories. Quotas set for each directory are called *subtree quotas*. When you link to the HCP system at the share level, you can limit the usage capacity per share based on the migration destination's hard quota for namespace capacity (namespace quota). For details on how to set quotas, see the *Administrator's Guide* or the *CLI Administrator's Guide*.

This section provides notes on quotas

- A CIFS client cannot view quota information from the Windows properties. To view quota information, use File Services Manager.
- For details on how to display the disk space for a CIFS client, see [Displaying disk capacity on page 8-53](#).
- A default quota cannot be set for a group.
- When the hard quota of the capacity for the namespace is larger than the hard limit on the quota for each file system and the quota per file system is exceeded, the write operation might fail even if the CIFS client seems to have free disk capacity.

Using CIFS access logs

System administrators and CIFS administrators can review CIFS access logs to monitor the access history of a CIFS share. System administrators must first configure the settings to determine whether and under what conditions to collect CIFS access logs.

CIFS access logs (`/var/log/cifs/log.CIFSaccess`) can be viewed on the **List of RAS Information** page (for **List of other log files**) in the **Check for Errors** dialog box. For details about how to view this page, see the *Administrator's Guide*. You can also save the most recent CIFS access log in a directory in the file system. For details about how to save CIFS access logs, see [Backing up the most recent CIFS access logs on page 3-17](#).

What to check before collecting CIFS access logs

Check the following items before collecting CIFS access logs:

- A system administrator must specify beforehand under what conditions the CIFS access logs are to be collected. The entire history of a CIFS client's accesses of a CIFS share is not logged; instead, the conditions under which CIFS access logs are collected differ depending on the individual settings of the CIFS service and the CIFS share.
- The conditions under which CIFS access logs are collected can be set up separately for each CIFS service or for each CIFS share. If the settings are configured for both the CIFS service and the CIFS share, priority is given to the CIFS share settings.
- The CIFS access log is output by a node unit to the same file. When the logged data exceeds the capacity that you set beforehand, the log file is

switched to the next generation. You can change capacity and the number of log files. For details, see the *Administrator's Guide*.

- Even if you set up the conditions under which CIFS access logs are collected, CIFS access logs might not be collected or might be collected at different conditions, depending on the access method from CIFS clients or the causes for why access attempts succeeded or failed.
- When the space occupied by log files on the OS disk reaches the limit, access history information is eliminated as follows:
 - If the settings are not configured to disable CIFS access log collection when the space occupied by log files reaches the limit, old log files are overwritten, and all access history recorded in overwritten log files is lost.
 - If the settings are configured to disable CIFS access log collection when the space occupied by log files reaches the limit, the collection of CIFS access logs is halted, and information on all subsequent accesses is lost.

Configure the settings for saving log files on the file system to prevent the loss of access history information. Log files saved on the file system use the following naming convention:

```
cifsaccesslog_node-host-name_YYYYMMDD_hhmmss.log
```

- To view a log file stored on a file system from a CIFS client, set up a CIFS share in the directory in which the file was stored and view the file using CIFS administrator privileges. To prevent improper access to stored log files by users, we recommend that you disable read and write access permissions on the CIFS share.
- When a log file saved on an OS disk is overwritten, an SNMP trap or email notification is sent. This might happen once every several minutes depending on the access state of the CIFS share. To prevent these kinds of notifications from being sent whenever a log file is overwritten, configure the settings to ensure that overwritten log files are saved on the file system.
- If a log file cannot be saved because there is not enough space on the destination file system, an SNMP trap or email notification is sent. This might happen once every several minutes depending on the access state of the CIFS share. To prevent these kinds of notifications from being sent whenever a log file cannot be saved, configure the settings to send a notification whenever file system usage exceeds the threshold value, and monitor usage of the file system.

Estimating log file size

When storing CIFS access log files, if the file system selected as the destination for stored log files runs out of space, the log files might not be stored. The system administrator must estimate the size of the log file that will be generated and set the capacity of the destination file system accordingly. In addition, the system administrator must periodically delete or transfer old log files to free disk space.

The table below shows the average size of log files generated over a period of one day with 1,000 clients accessing a CIFS share. The size of CIFS log files generated will depend on the networking environment and the status of CIFS client access utilization. The log file sizes presented in the following table can be taken as a rough benchmark, and administrators should allow sufficient leeway in estimating log file sizes.

Table 3-7 Sizes of CIFS access log files collected in an environment where 1,000 clients are accessing the system

CIFS access log configuration example	Trigger for collection of CIFS access log information	Log file size (MB/day)
Collect CIFS access log information related to connecting to or disconnecting from a CIFS share	Successful or unsuccessful in connecting to or disconnecting from a CIFS share	20
Collect CIFS access log information for any operation involving the writing of data	<ul style="list-style-type: none"> • Successful or unsuccessful in creating a file or writing data to a file • Successful or unsuccessful in creating a folder • Successful or unsuccessful in deleting a file or folder • Successful or unsuccessful in changing access permissions of a file or folder • Successful or unsuccessful in changing ownership of a file or folder • Successful or unsuccessful in connecting to or disconnecting from a CIFS share 	60
Collect CIFS access log information for any operations	<ul style="list-style-type: none"> • Successful or unsuccessful in displaying folder lists • Successful or unsuccessful in reading data • Successful or unsuccessful in creating a file or writing data to a file • Successful or unsuccessful in creating a folder • Successful or unsuccessful in deleting a file or folder • Successful or unsuccessful in reading the access permissions of a file or folder • Successful or unsuccessful in changing the access permissions of a file or folder • Successful or unsuccessful in changing the ownership of a file or folder • Successful or unsuccessful in connecting to or disconnecting from a CIFS share 	410

The log file sizes above are the result of 1,000 clients doing the following operations:

1. Each of the 1,000 clients uses the `dir` command to display 1,000 files in the CIFS share.
After these operations, the clients wait (do not perform any operations) for 5 minutes.
2. Each of the 1,000 clients copies one file in the CIFS share to another location in the same CIFS share.
After this operation, the clients wait (do not perform any operations) for 5 minutes.
3. The 1,000 clients repeat steps 1 and 2 (for the 1-day period during which the log data is collected).

Information written to the CIFS access log

Information is written to the CIFS access log in the following format:

```
date,time,process-id,user-name,IP-address-of-client-host,[CIFS-share-name],[result],[message-text],trigger,[details],[object-name]"
```

The information in the CIFS access log entries is described in the following table.

Table 3-8 Information in the CIFS access log entries

Item	Description
Date	Date on which event was generated in <code>YYYY/MM/DD</code> format.
Time	Time at which event was generated in <code>hh:mm:ss</code> format.
Process ID	Process ID of the process that generated the event.
User name	User name of the CIFS client who accessed the system.
IP address of client host	IP address of the CIFS client host.
CIFS share name	Name of the CIFS share. The name is not output if the CIFS share has already been deleted.
Result	Whether the system was successfully accessed. OK The system was successfully accessed. NG The system was not successfully accessed.
Message text	Message.
Trigger	The trigger that caused the entry to be written to the CIFS access log. <code>opendir</code> or <code>closedir</code> A list of folders was displayed.

Item	Description
	<p>open or close</p> <p>One of the following occurred:</p> <ul style="list-style-type: none"> - A data read was executed. - A file was created or a data write to the file was executed. - A folder was created. - A file or folder was deleted. - Access permissions for a file or folder were read. - Access permissions for a file or folder were modified. - The ownership on a file or folder was changed. - A file or folder was renamed. <p>mkdir</p> <p>A folder was created.</p> <p>unlink or rmdir</p> <p>A file or folder was deleted.</p> <p>sys_acl_get_file</p> <p>Access permissions of a file or folder were read.</p> <p>sys_acl_set_file</p> <p>Access permissions of a file or folder were modified.</p> <p>chown</p> <p>The ownership of a file or folder was changed.</p> <p>connect or disconnect</p> <p>A CIFS share was connected or disconnected.</p> <p>rename</p> <p>A file or folder name was renamed.</p> <p>set_owner</p> <p>One of the following was set as the owner of the file or folder:</p> <ul style="list-style-type: none"> - BUILTIN/Well-known SID account - SID unresolvable account <p>set_group</p> <p>One of the following was set in the file or folder group settings:</p> <ul style="list-style-type: none"> - BUILTIN/Well-known SID account - SID unresolvable account <p>set_dacl</p> <p>One of the following was set in the DACL settings for the file or folder:</p> <ul style="list-style-type: none"> - BUILTIN/Well-known SID account - SID unresolvable account <p>set_sacl</p> <p>An attempt was made to set an SACL for a file or folder.</p> <p>This will not be output if a user that lacks special permissions attempts to access a file or folder, resulting in an error.</p> <p>set_attr</p>

Item	Description
	<p>Unsupported attributes were included in the file attribute settings for a file or folder.</p>
<p>Details</p>	<p>Details of the event that triggered the collection of an entry to the CIFS access log.</p> <p>O_RDONLY When reading data from a file, the file was opened by the read-only attribute.</p> <p>O_WRONLY When creating a file or writing data to a file, the file was opened by the write-only attribute.</p> <p>O_RDWR When creating a file or writing data to a file, the file was opened by the read-write attribute.</p> <p>E When setting the file attribute for a file or folder, the Encrypted attribute was specified.</p> <p>C When setting the file attribute for a file or folder, the Compressed attribute was specified.</p> <p>O When setting the file attribute for a file or folder, the Offline attribute was specified.</p> <p>T When setting the file attribute for a file or folder, the Temporary attribute was specified.</p> <p>P When setting the file attribute for a file or folder, the SparseFile attribute was specified.</p> <p>L When setting the file attribute for a file or folder, the ReparsePoint attribute was specified.</p> <p>I When setting the file attribute for a file or folder, the NotContentIndexed attribute was specified.</p>
<p>Object name</p>	<p>If a file or folder was accessed, the name of the file or folder is output. If a CIFS share was connected to or disconnected from, the name of the CIFS share is output. If a file or folder was renamed, the absolute path both before and after the change is output in the following format: <i>old-absolute-path>new-absolute-path</i>.</p> <p>The name is not output if the target object has already been deleted.</p>

Backing up the most recent CIFS access logs

This subsection explains how to back up the most recent CIFS access logs.

Normally, when a CIFS access log file on the OS disk reaches its maximum size, it is automatically backed up to a pre-set directory in the file system. However, by using the `cifslogctl` command with the `--save` option, you can back up the most recent CIFS access log at any time, regardless of whether the log file has reached its maximum size.

For details on the trigger settings for collecting CIFS access logs, see the *Administrator's Guide* and the *CLI Administrator's Guide*. For details on how to specify a backup location for CIFS access logs, see the *CLI Administrator's Guide*.

Managing CIFS Client Users

This chapter explains how to manage CIFS client users.

- [Procedure for managing users](#)
- [Local user management](#)
- [User management in a domain](#)
- [Setting up an LDAP server for user mapping](#)
- [Manually registering a user ID and group ID](#)
- [User management when using the RFC 2307 schema](#)
- [Accessing CIFS shares when an HDI system is accessed from multiple domains](#)

Procedure for managing users

In an HDI system, you can manage the user information of file system users, such as UID, GID, and password, using the methods shown in the following table.

Table 4-1 User management methods supported by HDI

#	Item	Note
1	File Services Manager	To manage file system users with File Services Manager, you must register their user information.
2	NIS server	To manage file system users with an NIS server, you must register their user information.#
3	LDAP server (for user authentication)	To authenticate file system users with an LDAP server, you must register their user information.#
4	NT domain	To manage file system users that use an NT domain, you must perform one of the following operations: <ul style="list-style-type: none">• Register the user information with File Services Manager, an NIS server, or an LDAP server (for user authentication).• Set up user mapping.
5	Active Directory	To manage file system users that use Active Directory, you must perform one of the following operations: <ul style="list-style-type: none">• Register the user information with File Services Manager, an NIS server, or an LDAP server (for user authentication).• Set up user mapping.

#

For details on how to register NIS server information or user authentication LDAP server information into File Services Manager, see [Registering information in an NIS server or an LDAP server \(for user authentication\) on page 4-2](#).

Local user management

This section describes how to register, with File Services Manager, a user or group that has been registered with the NIS server or LDAP server (for user authentication). For details about how to register local users and group into File Services Manager, see the *Administrator's Guide*.

Registering information in an NIS server or an LDAP server (for user authentication)

When a user managed in the NIS server or LDAP server (for user authentication) uses local authentication to access a CIFS share in an HDI system, or uses the ACL function for a CIFS share in an HDI system, the user

registered with the NIS server or LDAP server (for user authentication) must also be registered to File Services Manager.

File Services Manager cannot authenticate a user who is registered with the NIS server or LDAP server (for user authentication). Therefore, to enable authentication of a user who is registered with the NIS server or LDAP server (for user authentication), a script is provided for registering a user from the NIS server or LDAP server (for user authentication) to File Services Manager.

Also, if you want to use a group managed by an NIS server or LDAP server (for user authentication) to access a CIFS share in an HDI system, you must map the group registered on the NIS server or LDAP server (for user authentication) with a group registered to File Services Manager.

File Services Manager cannot use a group registered on the NIS server or LDAP server (for user authentication) in an HDI system. Therefore, to use a group that is registered on the NIS server or LDAP server (for user authentication) to access a CIFS share in an HDI system, a script is provided for mapping a group that is registered in File Services Manager to a group that is registered on the NIS server or LDAP server (for user authentication).

Overview of functionality

The following shows the procedure for adding or deleting users, and the procedure for registering or removing group mappings. A CSV file containing user information is used to add or delete users. Likewise, a CSV file containing mapping information is used to register or remove group mappings. The formats of these CSV files are shown after the procedures. For details on the format of the CSV files above, see [CSV file formats on page 4-4](#).

(a) Procedure for adding or deleting users

If a user managed in the NIS server or LDAP server (for user authentication) accesses a CIFS share in an HDI system, you must register the user to File Services Manager. This subsection describes how to register a user registered on the NIS server or LDAP server (for user authentication) to File Services Manager.

- The users registered here are used when the CIFS file share is accessed.
 - The password is used during local authentication, when the CIFS file share is accessed.
1. Create a CIFS file share (limiting the clients from which it can be accessed).
In step 2, you must save a non-encrypted password to a file. Therefore, to prevent other users from seeing the password, we strongly recommend that you restrict access to the CIFS share to be created.
 2. Save the CSV file in the directory created in the step 1.
Perform a virus scan on the CSV file to be saved to confirm that there is no problem.
 3. Use SSH to log on to the HDI node.

4. Execute the script `sudo cifsusredit`, which is used for registering, deleting, or listing users.
5. Log off from the HDI node.
6. Delete the CSV file created in step 2 and the shared directory created in step 1.
7. Perform steps 1 to 6 in the same way for all other nodes comprising the cluster.

(b) Procedure for registering or removing a group mapping

If you want to use a group managed by an NIS server or LDAP server (for user authentication) to access a CIFS share in an HDI system, you must map the group with a group registered to File Services Manager. This subsection describes how to map a group registered on the NIS server or LDAP server (for user authentication) to File Services Manager. Note that the mapping registered here is used for the ACL of the CIFS file share resource.

1. Create a CIFS file share.
We strongly recommend that you restrict access to the CIFS share to be created.
2. Save the CSV file in the directory created in the step 1.
Perform a virus scan on the CSV file to be saved.
3. Use SSH to log on to the HDI node.
4. Execute the script `sudo cifsgrpedit`, which is for registering, deleting, or listing group mappings.
5. Log off from the HDI node.
6. Delete the CSV file created in step 2.
7. Perform steps 1 to 6 in the same way for all other nodes comprising the cluster.

CSV file formats

The data files are in CSV (comma-separated value) format.

The CSV file formats are as follows:

- Commas (,) delimit fields, with no space characters before or after each field. All space characters are interpreted as part of the field's value.

Example: Specifying CSV file entries

```
field-1-1,field-1-2  
field-2-1,field-2-2  
field-3-1,field-3-2
```

- When the value of a field contains double quotation marks ("), specify two quotation marks for each one, and then enclose the entire field in quotation marks.

Example: Specifying field1,field"2"

```
field1,"field"2""
```

- Enclose in double quotation marks (") fields whose value contains a comma (,).

Example: field1,field,2

```
field1,"field,2"
```

- Insert a line-feed character after each line.

(a) Format for user registration files

The format for user registration files is as follows. Since information for only one user can be specified on each line, use multiple lines to specify multiple users.

```
user-name,password  
user-name,password  
user-name,password
```

(b) Format for group mapping

The format for group mapping files is as follows. Since information for only one group can be specified on each line, use multiple lines to specify multiple group mappings.

```
NIS-server-or-other-external-group-name,group-name-registered-with-File-  
Services-Manager  
NIS-server-or-other-external-group-name,group-name-registered-with-File-  
Services-Manager  
NIS-server-or-other-external-group-name,group-name-registered-with-File-  
Services-Manager
```

- Characters that cannot be specified for group mappings
Operation is not guaranteed when the following characters are used:
\
/[] : | < > + = ; , ? * "

Script for registering, deleting, or listing users

This subsection describes the script that is used for registering, deleting, or listing users.

Name

```
cifsusredit
```

Syntax

```
sudo cifsusredit option [csv-file]
```

Description

This command registers, deletes, or lists CIFS user.

Arguments

option

Specify `add`, `delete`, or `list` to perform the following operations. This argument is required.

- `add`

Registers users indicated in the specified *csv-file* to File Services Manager. When `add` is specified for *option*, the *csv-file* argument is required. The execution results are output to the standard output.

- `delete`

Deletes users indicated in the specified *csv-file* from File Services Manager. When `delete` is specified for *option*, the *csv-file* argument is required. The execution results are output to the standard output.

- `list`

Outputs user names registered with File Services Manager in the standard output.

csv-file

Specify a CSV file that contains user information.

Return value

When registration or deletion of all users specified in the CSV file terminated normally, 0 is returned. When abnormally terminated, a non-zero value is returned.

Notes:

- *csv-file*: the name of the CSV file saved in the CIFS file share directory. If the CSV file `file.csv` is saved in the shared directory `/mnt/test1/test1`, specify `/mnt/test1/test1/file.csv` for this command argument. This argument must be specified when `add` or `delete` is specified for the *option* argument.
- When saving a CSV file with user names and passwords on a CIFS file share directory, be sure to restrict access to the shared directory, so that other clients cannot access it.
- Once this command is executed, be sure to delete the CSV file with the user names and passwords as soon as possible.
- For the user names specified in the CSV file, specify the user names registered in the NIS server or LDAP server (for user authentication).
- If a specified CSV file contains a user name that has already been registered with File Services Manager, the password specified in the CSV file will overwrite the current password.
- The `LF` or `CR+LF` line return codes can be used.
- Do not specify two-byte codes. Operation is not guaranteed when two-byte codes are specified.

Script for CIFS group mapping

This subsection describes the script that is used for mapping groups.

Name

`cifsgrpedit`

Syntax

```
sudo cifsgrpedit option [csv-file]
```

Description

This command enables a group registered with the NIS server or LDAP server (for user authentication) to be used as an HDI system group.

Arguments

option

Specify `add`, `delete`, or `list` to perform the following operations. This argument is required.

- `add`

Registers mapping information indicated in the specified *csv-file* to File Services Manager. When `add` is specified for *option*, the *csv-file* argument is required. The execution results are output to the standard output.

- `delete`

Deletes mapping information indicated in the specified *csv-file* from File Services Manager. When `delete` is specified for *option*, the *csv-file* argument is required. The execution results are output to the standard output.

- `list`

Outputs group names registered with File Services Manager in the standard output.

csv-file

Specify a CSV file that contains group mapping information.

Return value

This command returns the value `0`, if all the groups specified in the CSV file are successfully mapped. If the mapping ends abnormally, this command returns a value other than `0`.

Note:

- *csv-file* is the name of the CSV file saved in the CIFS file share directory. If the CSV file `file.csv` is saved in the shared directory `/mnt/test1/test1`, specify `/mnt/test1/test1/file.csv` for this command argument. This argument must be specified when `add` or `delete` is specified for the *option* argument.
- For the group names specified in the CSV file, specify the group names registered in the NIS server or LDAP server (for user authentication).
- If a specified CSV file contains a group name that has already been registered with File Services Manager, group mapping will fail.
- The `LF` or `CR+LF` codes can be used as line breaks.

- Do not specify two-byte codes. Operation is not guaranteed when two-byte codes are specified.

Note on users managed by an NIS server or an LDAP server (for user authentication)

When a user managed by an NIS server or an LDAP server (for user authentication) accesses a CIFS share in an HDI system or uses the ACL function on a CIFS share in an HDI system, the comment set when the user was registered is used for the display of the ACL.

Notes on local users and group registration

The following notes are about registering local users and groups.

- A user who performs CIFS access can belong to a maximum of 1,023 groups, including the primary group. If the user belongs to more than 1,023 groups, use user mapping.
- If local users and groups with the same names as built-in Windows users and groups are registered in File Services Manager, they might be recognized as the built-in Windows users and groups. In such cases, those local users and groups cannot be set as ACLs or owners from a CIFS client.

User management in a domain

If you use user mapping, note the following:

- The HDI system does not recognize built-in Windows users and groups.
- Nested groups of Windows are valid on the HDI system when the Active Directory domain is in native mode.

If you do not use user mapping, note the following:

- The valid group for a user is the one registered with File Services Manager, NIS, and so on. Groups on the domain controller are invalid.

Setting up an LDAP server for user mapping

When **Use user mapping using LDAP.** is selected for the user mapping method, an LDAP server must be set up to perform user mapping. [Table 4-2 Support status of the LDAP server for user mapping in the HDI system on page 4-9](#) lists the support status of the LDAP server for user mapping in the HDI system.

This section provides notes and example settings for when you use OpenLDAP, ADAM, or Sun Java System Directory Server to create an LDAP server for user mapping.

Table 4-2 Support status of the LDAP server for user mapping in the HDI system

LDAP server		Status of support
Open LDAP	Linux	Supported
	Solaris	Supported
ADAM		Supported
Sun Java System Directory Server		Supported

Precautions when setting up an LDAP server

When an LDAP server is initialized or set up again, the CIFS service needs to be restarted. Before restarting the CIFS service, make sure no users are accessing a CIFS file share.

In addition, after restarting the CIFS service, delete any user mapping information from the cache of the CIFS service environment.

Precautions when using OpenLDAP to set up an LDAP server

This subsection explains what precautions to take when using OpenLDAP to set up an LDAP server.

If you use OpenLDAP to set up an LDAP server, the user mapping functionality cannot be used with an OpenLDAP server using transport layer security (TLS), a protocol for sending and receiving encrypted information over the Internet.

The maximum search number (number of entries returned for a search request from an LDAP client) can be specified for an OpenLDAP LDAP server.

- The default is 500 entries.
- If the entries of user information and user mapping information contained in an LDAP server exceed the maximum number, an attempt to download user mapping information might fail in the **List of RAS Information** page (for *Batch-download*) of the **Check for Errors** dialog box, or you might not be able to list information on such pages as the **List of Quota Information** page of the **Edit Quota** dialog box. Also, in the **Create and Share File System, Add Share, or Edit Share** dialog box, the **Access Control** tab does not correctly display **All Users** and **All Groups of Special permitted users/groups**. Therefore, add the following `sizelimit` directive to the LDAP server definitions:

```
sizelimit -1
```

Precautions when using ADAM to set up an LDAP server

The following explains precautions for using ADAM to set up an LDAP server.

The maximum search number (number of entries returned for a search request from an LDAP client) can be specified for an ADAM LDAP server.

- The default is 1,000 entries.
- If the entries of user mapping information in an LDAP server exceed the maximum number, an attempt to download user mapping information fails on the **List of RAS Information** page (for `Batch-download`) of the **Check for Errors** dialog box. Therefore, increase the `MaxPageSize` limit so that the maximum number of search results exceeds the sum of the number of users and number of groups managed.

The following shows the procedure to increasing the `MaxPageSize` limit. For details about the **Edit ADAM ADSI** tool and the terms used in this procedure, see the appropriate Microsoft documentation.

1. Use the **Edit ADAM ADSI** tool to connect to the configuration partition.
2. Expand the console tree, and click **CN=Services, CN=Windows NT, CN=Directory Service**, and **CN=Query-Policies**, in order.
3. In the details window, double-click **CN=Default Query Policy**, and then in the properties window, double-click the **IDAPAdminLimits** attribute. Edit the attribute value.
4. Select **MaxPageSize=1000**, and then click the **Delete** button.
5. Enter `MaxPageSize=maximum-number`, and then click the **Add** button. For *maximum-number*, use the sum of the maximum users plus the maximum groups, based on the user ID scope and the group ID scope that is specified when configuring user mappings for File Services Manager.
6. Click the **OK** button twice to complete the settings.

Precautions when using Sun Java System Directory Server to set up an LDAP server

The following explains precautions for using Sun Java System Directory Server to set up an LDAP server.

The maximum search number (number of entries returned for a search request from an LDAP client) can be specified for a Sun Java System Directory Server LDAP server.

- The default is 2,000 entries.
- If the entries of user information and user mapping information contained in an LDAP server exceed the maximum number, an attempt to download user mapping information might fail on the **List of RAS Information** page (for `Batch-download`) of the **Check for Errors** dialog box, or you might not be able to list the information on such pages as the **List of Quota Information** page of the **Edit Quota** dialog box. Also, in the **Create and Share File System, Add Share**, or **Edit Share** dialog box, the **Access Control** tab does not correctly display **All Users** and **All Groups** of **Special permitted users/groups**. Therefore, you must

change the maximum number of search results to **No limit** for the LDAP server created by using Sun Java System Directory Server.

The description of how to change the maximum number of search results to **No limit** is provided below. For details on the terminology used in this procedure, see Sun Java System Directory Server documentation.

1. In the **Configuration** tab of the console for the LDAP server set up by Sun Java System Directory Server, display the directory tree, and choose **Performance**.
2. In the right-hand panel, choose the **Client control** tab.
3. Select the **Unlimited** check boxes for **LDAP size limit** and **Search right**.
4. Click the **Save** button.
A message is displayed prompting you to restart Sun Java System Directory Server.
5. Click the **OK** button.
6. Click the **Tasks** tab, and click the restart button for Sun Java System Directory Server.
A dialog box is displayed to confirm the restart operation. Click **Yes**.
7. Click the **Close** button to close the **Restart Directory Server** dialog box.

Examples settings for when using OpenLDAP to set up an LDAP server

This section contains example settings for when using OpenLDAP to set up an LDAP server.

Creating a schema file

To use LDAP user mapping, create a schema file for defining attributes and object classes that will be recognized by an LDAP server created in OpenLDAP. These attributes and object classes need to be defined to store the user IDs and group IDs converted by the user mapping.

The HDI system provides a schema file, `samba.schema`, required for user mapping. Obtain the schema file from the following directory by using the `scp` command from the remote host:

```
/usr/share/doc/cifs/examples/samba.schema
```

When creating a schema file for an LDAP server set up using OpenLDAP, define the following attributes and object classes.

```
attributetype ( 1.3.6.1.4.1.7165.2.1.20 NAME 'sambaSID'
    DESC 'Security ID'
    EQUALITY caseIgnoreIA5Match
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.26{64} SINGLE-VALUE )
objectclass ( .3.6.1.4.1.7165.2.2.7 NAME 'sambaUnixIdPool' SUP top AUXILIARY
    DESC 'Pool for allocating UNIX uids/gids'
    MUST ( uidNumber $ gidNumber ) )
objectclass ( 1.3.6.1.4.1.7165.2.2.8 NAME 'sambaIdmapEntry' SUP top AUXILIARY
```

```
DESC 'Mapping from a SID to an ID'  
MUST ( sambaSID )  
MAY ( uidNumber $ gidNumber ) )  
objectclass ( 1.3.6.1.4.1.7165.2.2.9 NAME 'sambaSidEntry' SUP top STRUCTURAL  
DESC 'Structural Class for a SID'  
MUST ( sambaSID ) )
```

Once the schema file has been created or obtained, to load the schema file for using user mapping, add the `include` directive to the LDAP server definition.

The following is an example of how to use the `include` directive for when the schema file is in the `/etc/ldap/schema` directory:

```
include /etc/ldap/schema/samba.schema
```

Setting the index directive

Since search performance for an LDAP server set up by OpenLDAP may degrade when the number of user IDs and group IDs stored on the LDAP server increases, set the `index` directive. When user mapping is used, specify the `index` directive as follows in the LDAP server definition:

```
index uidNumber,gidNumber,objectClass,sambaSID eq
```

- When the `index` directive is changed, the index needs to be recreated based on the current contents of the LDAP server database. Use the `slapindex` command provided by OpenLDAP to recreate the index.
- When executing the `slapindex` command, stop the LDAP server momentarily, and then restart it after the `slapindex` command is executed.

Example settings for when using ADAM to set up an LDAP server

This section contains example settings for when using ADAM to set up an LDAP server.

Creating a schema file

To use LDAP user mapping, create a schema file for defining attributes and object classes that is recognized by an LDAP server created in ADAM. With an LDAP server, these attributes and object classes need to be defined to store the user IDs and group IDs converted by the user mapping.

The HDI system provides a schema file, `samba.ldf`, required for LDAP user mapping. Obtain the schema file from the following directory by using the SCP functionality from the remote host:

```
/usr/share/doc/cifs/examples/samba.ldf
```

When creating a schema file for an LDAP server set up using ADAM, define the following attributes and object classes.

```

dn: CN=uidNumber,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: top
objectClass: attributeSchema
cn: uidNumber
attributeID: 1.3.6.1.1.1.1.0
attributeSyntax: 2.5.5.9
isSingleValued: TRUE
showInAdvancedViewOnly: FALSE
adminDisplayName: uidNumber
adminDescription: An integer uniquely identifying a user in an
  administrative domain
oMSyntax: 2
searchFlags: 1
LDAPDisplayName: uidNumber
systemOnly: FALSE
systemFlags: 16

dn: CN=gidNumber,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: top
objectClass: attributeSchema
cn: gidNumber
instanceType: 4
attributeID: 1.3.6.1.1.1.1.1
attributeSyntax: 2.5.5.9
isSingleValued: TRUE
showInAdvancedViewOnly: FALSE
adminDisplayName: gidNumber
adminDescription: An integer uniquely identifying a group in an
  administrative domain
oMSyntax: 2
searchFlags: 1
LDAPDisplayName: gidNumber
systemOnly: FALSE
systemFlags: 16

dn: CN=sambaSID,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: top
objectClass: attributeSchema
cn: sambaSID
instanceType: 4
attributeID: 1.3.6.1.4.1.7165.2.1.20
attributeSyntax: 2.5.5.12
isSingleValued: TRUE
showInAdvancedViewOnly: FALSE
adminDisplayName: sambaSID
adminDescription: Security ID
oMSyntax: 64
searchFlags: 1
LDAPDisplayName: sambaSID
systemOnly: FALSE
systemFlags: 16

dn: CN=sambaUnixIdPool,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: top
objectClass: classSchema
cn: sambaUnixIdPool
instanceType: 4
possSuperiors: container
subClassOf: top
governsID: 1.3.6.1.4.1.7165.1.2.2.7
rDNAttID: cn
showInAdvancedViewOnly: FALSE
adminDisplayName: sambaUnixIdPool

```

```

adminDescription: Pool for allocating UNIX uids/gids
objectClassCategory: 3
LDAPDisplayName: sambaUnixIdPool
systemOnly: FALSE
systemPossSuperiors: organizationalUnit
systemPossSuperiors: domainDNS
mustContain: uidNumber
mustContain: gidNumber
defaultSecurityDescriptor:
  D: (A;;;RPWPCRCCDCLCLOCRCWOWSDDTSW;;;DA) (A;;;RPWPCRCCDCLCLOCRCWOWSD
  DDTSW;;;SY) (A;;;RPLCLORC;;;AU)
systemFlags: 16
defaultHidingValue: FALSE

dn: CN=sambaIdmapEntry,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: top
objectClass: classSchema
cn: sambaIdmapEntry
instanceType: 4
possSuperiors: container
subClassOf: top
governsID: 1.3.6.1.4.1.7165.1.2.2.8
rDNAttID: cn
showInAdvancedViewOnly: FALSE
adminDisplayName: sambaIdmapEntry
adminDescription: Mapping from a SID to an ID
objectClassCategory: 3
LDAPDisplayName: sambaIdmapEntry
systemOnly: FALSE
systemPossSuperiors: organizationalUnit
systemPossSuperiors: domainDNS
mustContain: sambaSID
mayContain: gidNumber
mayContain: uidNumber
defaultSecurityDescriptor:
  D: (A;;;RPWPCRCCDCLCLOCRCWOWSDDTSW;;;DA) (A;;;RPWPCRCCDCLCLOCRCWOWSD
  DDTSW;;;SY) (A;;;RPLCLORC;;;AU)
systemFlags: 16
defaultHidingValue: FALSE

dn: CN=sambaSidEntry,CN=Schema,CN=Configuration,DC=X
changetype: ntdsSchemaAdd
objectClass: top
objectClass: classSchema
cn: sambaSidEntry
instanceType: 4
possSuperiors: container
subClassOf: top
governsID: 1.3.6.1.4.1.7165.1.2.2.9
rDNAttID: sambaSID
showInAdvancedViewOnly: FALSE
adminDisplayName: sambaSidEntry
adminDescription: Structural Class for a SID
objectClassCategory: 1
LDAPDisplayName: sambaSidEntry
systemOnly: FALSE
systemPossSuperiors: organizationalUnit
systemPossSuperiors: domainDNS
mustContain: sambaSID
defaultSecurityDescriptor:
  D: (A;;;RPWPCRCCDCLCLOCRCWOWSDDTSW;;;DA) (A;;;RPWPCRCCDCLCLOCRCWOWSD
  DDTSW;;;SY) (A;;;RPLCLORC;;;AU)
systemFlags: 16
defaultHidingValue: FALSE

```

Once the schema file has been created or obtained, load the schema file for using user mapping by entering as one line and then executing the following command in the command prompt:

```
ldifde -i -f C:\samba.ldf -s localhost:port-number -j . -k -c
"CN=Schema,CN=Configuration,DC=X" #schemaNamingContext
```

In this example, the schema file is saved as `C:\samba.ldf`. For *port-number*, specify the LDAP port number specified when ADAM was installed. The `ldifde` command exists on the system when ADAM or Active Directory is installed. To use the ADAM `ldifde` command, from the **Start** menu, choose **All Programs**, then **ADAM**, and then **ADAM Tool Command Prompt**.

Setting index

Since search performance for an LDAP server set up by ADAM may degrade when the number of user IDs and group IDs stored on the LDAP server increases, set `index`.

When ADAM is used and the schema is extended, `index` is set for the extended `uidNumber`, `gidNumber`, and `sambaSID`. The following shows the procedure to set `index` for the existing system attribute `objectClass`. For details about the **Edit ADAM ADSI** tool and the terms used in this procedure, see the appropriate Microsoft documentation.

1. Use the **Edit ADAM ADSI** tool to connect to the scheme partition.
2. Expand the console tree, and in the Details window, double-click **cn=Object-Class**.
3. In the Properties window, double-click the **searchFlags** attribute, and then edit the attribute value.

Since the value set is 8, change it to 9. If it has already been changed to another value, specify the following based on the value set:

- If an odd number is set, leave the value as set.
 - If an even number is set, increase the set value by 1.
4. Click the **OK** button twice to close the dialog boxes.

Example settings for when using Sun Java System Directory Server to set up an LDAP server

This section contains example settings for when using Sun Java System Directory Server to set up an LDAP server.

Creating a schema file

To use LDAP user mapping, create a schema file for defining attributes and object classes that is recognized by an LDAP server created in Sun Java System Directory Server. With an LDAP server, these attributes and object classes need to be defined to store the user IDs and group IDs converted by the user mapping.

The HDI system provides a schema file, `samba.ldif`, required for LDAP user mapping. Obtain the schema file from the following directory by using the `scp` command from the remote host:

```
/usr/share/doc/cifs/examples/samba.ldif
```

When creating a schema file for an LDAP server set up using Sun Java System Directory Server, define the following attributes and object classes.

```
dn: cn=schema
changetype:modify
add: attributeTypes
attributeTypes: ( 1.3.6.1.4.1.7165.2.1.20 NAME 'sambaSID' DESC 'Security ID'
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 SINGLE-VALUE X-ORIGIN 'user defined' )
-
add: objectClasses
objectClasses: ( 1.3.6.1.4.1.7165.1.2.2.7 NAME 'sambaUnixIdPool' SUP top
  AUXILIARY MUST ( uidNumber $ gidNumber ) X-ORIGIN 'user defined' )
-
add: objectClasses
objectClasses: ( 1.3.6.1.4.1.7165.1.2.2.8 NAME 'sambaIdmapEntry' SUP top
  AUXILIARY MUST sambaSID MAY ( uidNumber $ gidNumber ) X-ORIGIN
  'user defined' )
-
add: objectClasses
objectClasses: ( 1.3.6.1.4.1.7165.1.2.2.9 NAME 'sambaSidEntry' SUP top
  STRUCTURAL MUST sambaSID X-ORIGIN 'user defined' )
-
```

Once the schema file has been created or obtained, to load the schema file for using user mapping, execute the following command to extend the schema. If a password is requested, enter that set for `cn=Directory Manager`, during installation.

```
#ldapmodify -h host-name -p port-number -D "cn=Directory Manager" -w - -f
samba.ldif
```

Use the `ldapmodify` command provided by Sun Java System Directory Server (and not that provided by OpenLDAP). For *host-name*, specify the host name of the LDAP server created by using Sun Java System Directory Server. For *port-number*, specify the LDAP port number specified when Sun Java System Directory Server was installed.

Setting index

Since search performance for an LDAP server set up by Sun Java System Directory Server may degrade when the number of user IDs and group IDs stored on the LDAP server increases, set `index`.

When Sun Java System Directory Server is used, an equality index is set for the `objectClass` attribute. When user mapping is used, we recommend that you set an equality index (`eq`) for `uidNumber`, `gidNumber`, and `sambaSID`, in the Sun Java System Directory Server definition.

The following shows the procedure for setting an equality index (`eq`) for `uidNumber`, `gidNumber`, and `sambaSID`. For details about the terms used in

this procedure, see the appropriate Sun Java System Directory Server documentation.

1. In the **Configuration** tab of the console for the LDAP server set up by Sun Java System Directory Server, expand the **Data** node, and select a suffix for which to generate an index.
2. In the right-hand panel, choose the **Indexes** tab.
The system index table cannot be changed.
3. Add an index for the attributes of the **Additional Index** table.
4. To add an index to an attribute for which no index exists, click the **Add attribute** button.
A dialog box is displayed, so select `uidNumber`, `gidNumber`, `sambaSID` for which to generate an index, and click **OK**.
5. To change an attribute index, in the **Additional Index** table, select the type check box for the index maintained for the attribute.
Check that the **Equality** index check box is selected for `uidNumber`, `gidNumber`, and `sambaSID`, and clear the **Presence** index check box. Do not select any other check boxes.
6. Click the **Save** button to save the new index settings.
A dialog box is displayed to indicate that an update of the database file is required to use a new index.
The suffix index can be regenerated or re-initialized. Since mapping information has yet to be registered, select **Do nothing**.

Manually registering a user ID and group ID

This section describes how to register your desired user IDs and group IDs manually when using user mapping.

How to register IDs with Active Directory

When you select **Use user mapping using Active Directory schema** as a user mapping method, you must manually register your desired user IDs and group IDs in the user management window of Active Directory.

This subsection describes how to do this.

Registering a group ID

To manually register a group ID:

1. In the **Active Directory Users and Computers** window of the domain controller, open the Properties window of the target group.
2. Select the **UNIX Attributes** tab.
3. Select an appropriate item from the **NIS Domain** pull-down menu.
4. In the **GID (Group ID)** text box, change the value to your desired group ID.
5. Click the **Apply** button.

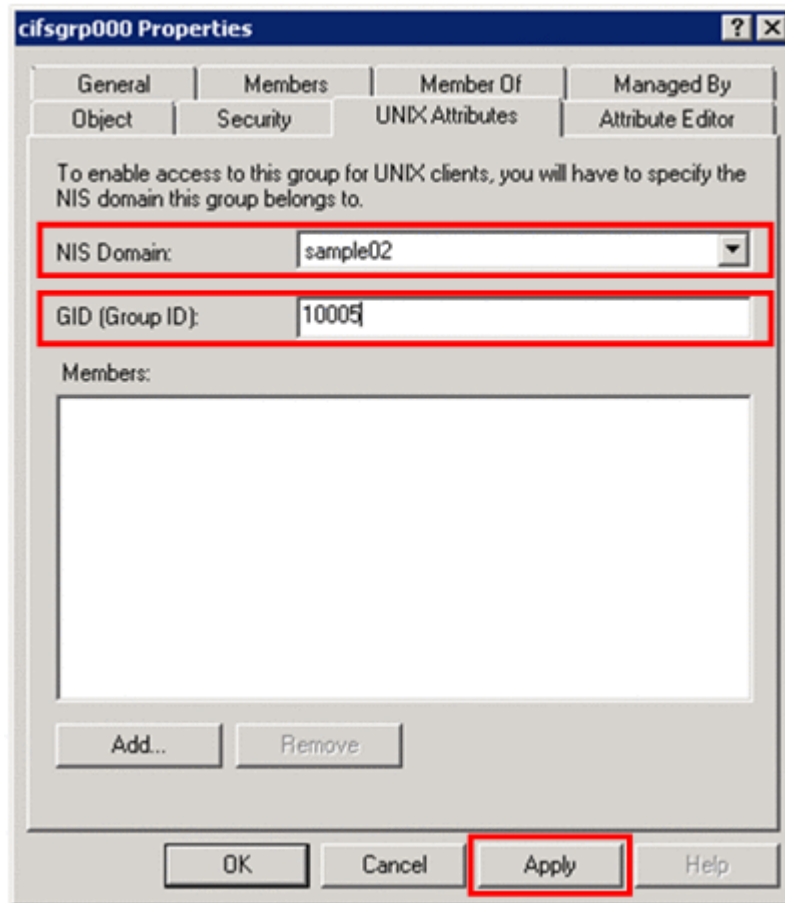


Figure 4-1 Below is an example of a window displaying the UNIX Attributes tab on the Properties page of the group

Registering a user ID

To manually register a user ID:

1. In the **Active Directory Users and Computers** window of the domain controller, open the Properties window of the target user.
2. In the **Member Of** tab, confirm that the primary group has a UNIX attribute GID.
3. Select the **UNIX Attributes** tab.
4. Select an appropriate item from the **NIS Domain** pull-down menu.
5. In the **UID** text box, change the value to your desired user ID.
6. From the **Primary group name/GID** pull-down menu, select an appropriate primary group.
7. Click the **Apply** button.

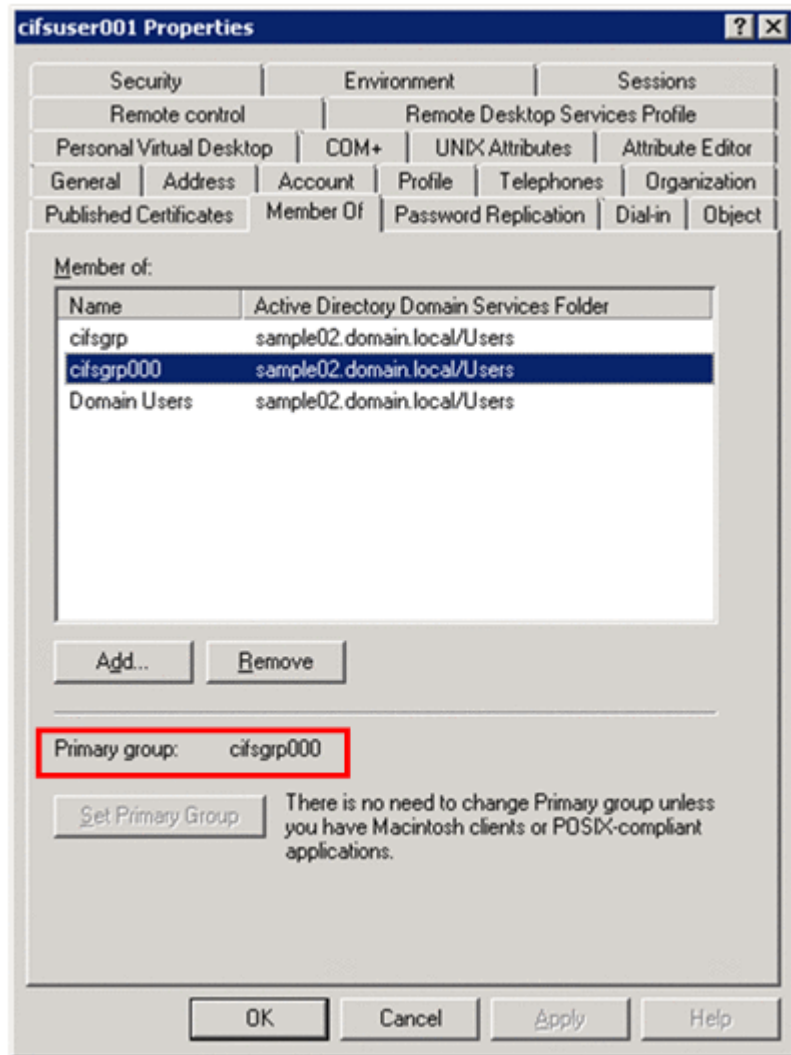


Figure 4-2 Below is an example of a window displaying the Member Of tab on the Properties page of the user




Figure 4-3 Below is an example of a window displaying the UNIX Attributes tab on the Properties page of the user

How to register IDs with an LDAP server

When you select **Use user mapping using LDAP.** on the **CIFS Service Management** page (**Setting Type:** *User mapping*) as a user mapping method and also select **Allocate manually**, you must manually register your desired user IDs and group IDs with an LDAP server.

This subsection describes how to do the procedure described above.

Note:

After you manually register IDs with an LDAP server, do not change the ID allocation method from **Allocate manually** to **Allocate automatically**. Doing this might duplicate user mapping information.

Registering a group ID

To manually register a group ID:

1. On the LDAP server, prepare a file that contains information of the target group in the following format:

```
dn: sambaSID=SID-of-a-group-in-Active-Directory-or-the-NT-domain,DN-of-user-mapping-LDAP
objectClass: sambaIdmapEntry
objectClass: sambaSidEntry
gidNumber: UNIX-attribute-GID-to-be-allocated-to-the-group
sambaSID: SID-of-a-group-in-Active-Directory-or-the-NT-domain
```

Example:

```
dn: sambaSID=S-1-5-21-848980995-581375927-1041525310-53490,dc=test,dc=local
objectClass: sambaIdmapEntry
objectClass: sambaSidEntry
gidNumber: 200000
sambaSID: S-1-5-21-848980995-581375927-1041525310-53490
```

2. Execute the `ldapadd` command in the following format:

```
ldapadd -f file-name-that-contains-group-information -x -D "common-name-of-the-LDAP-administrator, DN-of-user-mapping-LDAP" -w password-of-the-LDAP-administrator
```

Example:

```
ldapadd -f entries.ldif -x -D "cn=Manager,dc=test,dc=local" -w adminpass
```

Registering a user ID

To register a user ID:

1. On the LDAP server, prepare a file that contains information of the target user in the following format:

```
dn: sambaSID=SID-of-a-user-in-Active-Directory-or-the-NT-domain,DN-of-user-mapping-LDAP
objectClass: sambaIdmapEntry
objectClass: sambaSidEntry
uidNumber: UNIX-attribute-UID-to-be-allocated-to-the-user
sambaSID: SID-of-a-user-in-Active-Directory-or-the-NT-domain
```

Example:

```
dn: sambaSID=S-1-5-21-848980995-581375927-1041525310-53491,dc=test,dc=local
objectClass: sambaIdmapEntry
objectClass: sambaSidEntry
uidNumber: 200001
sambaSID: S-1-5-21-848980995-581375927-1041525310-53491
```

2. Execute the `ldapadd` command in the following format:

```
ldapadd -f file-name-that-contains-user-information -x -D "common-name-of-the-LDAP-administrator, DN-of-user-mapping-LDAP" -w password-of-the-LDAP-administrator
```

Example:

```
ldapadd -f entries.ldif -x -D "cn=Manager,dc=test,dc=local" -w adminpass
```

How to delete IDs registered with an LDAP server

To delete a user ID or group ID that is registered with an LDAP server:

1. On the LDAP server, execute the `ldapdelete` command in the following format:

```
ldapdelete -x -D "common-name-of-the-LDAP-administrator, DN-of-user-mapping-LDAP" "sambaSID=SID-of-a-user-in-Active-Directory-or-the-NT-domain, organization-unit-name-of-user-mapping-LDAP, DN-of-user-mapping-LDAP" -w password-of-the-LDAP-administrator
```

Example:

```
ldapdelete -x -D "cn=Manager,dc=test,dc=local" "sambaSID=S-1-5-21-848980995-581375927-1041525310-53491,ou=idmap,dc=test,dc=local" -w adminpass
```

2. In File Services Manager, from the **Access Protocol Configuration** dialog box, on the **CIFS Service Maintenance** page, click the **Clear User Map Cache File** button to delete the cache file.

User management when using the RFC 2307 schema

This section serves as a supplement for when the Active Directory Schema is used for the user mapping method and users are managed by specifying **Using LDAP as a network information service (RFC2307)** for **Name service switch** under **User mapping setup** on the **CIFS Service Management** page (**Setting Type**: User mapping).

In an HDI system, from among the UNIX attribute values of domain users to be used when a CIFS client accesses a CIFS share, a UNIX attribute value to be used for the group ID of a primary group can be selected from either of the following two values:

- `gidNumber` value for the group (indicated by the `primaryGroupID` of a UNIX attribute) that the user belongs to

This corresponds to the group (Domain Users in the example in the following figure) that is displayed for **Primary group** below in the **Member Of** tab of the Properties window for the Active Directory user.

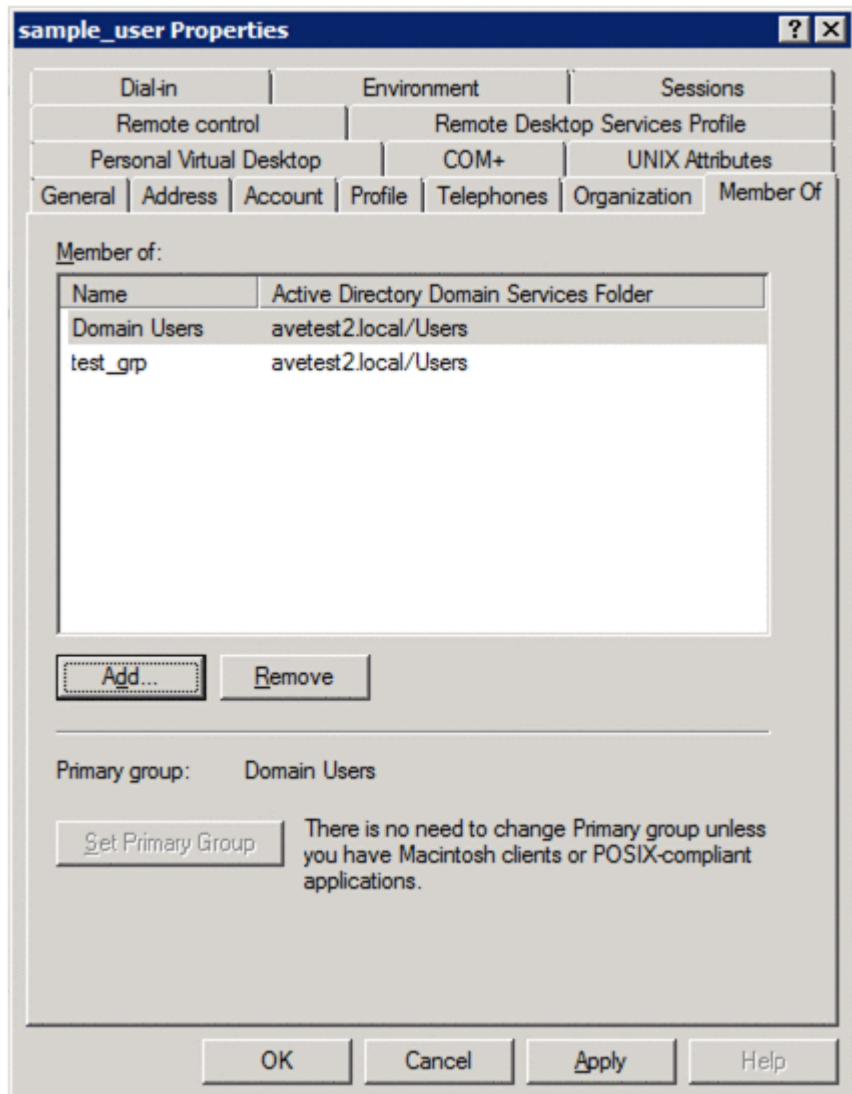


Figure 4-4 Window that displays the Member Of tab

- `gidNumber` value for the user
This corresponds to a group (`sample_group` in the example in the following figure) that is displayed for **Primary group name/GID** below in the **UNIX Attributes** tab of the Properties window for the Active Directory user.

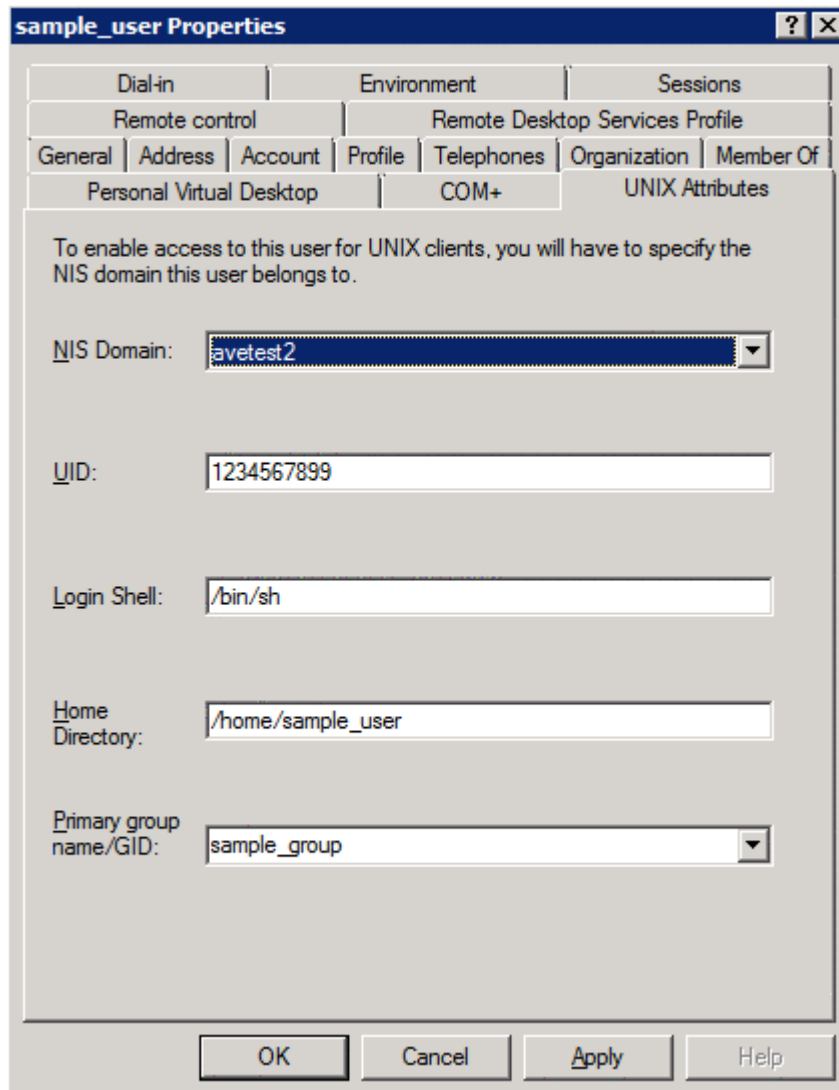


Figure 4-5 Window that displays the UNIX Attributes tab

HDI serves as a group for CIFS clients to access CIFS shares. By default, HDI runs by using `gidNumber` values for groups. To operate HDI by using `gidNumber` values of users, the `use_gidnumber` option must be specified to run the `cifsopstset` command, and the CIFS service settings must be changed.

Accessing CIFS shares when an HDI system is accessed from multiple domains

Users who belong to a domain that is in a trust relationship with the domain that the nodes belong to can also access the CIFS shares in an HDI system. For details about domains that allow access to an HDI system, see the *Installation and Configuration Guide*. When the domain configuration is changed, it might take time to access the CIFS shares. After you receive

contact from the domain administrator regarding the changes to the domain configuration, change the configuration definition of the CIFS service to fit these changes. For details about how to change the configuration definition of the CIFS service, see the *Administrator's Guide*.

User Authentication for CIFS Clients

This chapter contains notes regarding user authentication for CIFS clients.

- [Local authentication](#)
- [NT domain authentication](#)
- [Active Directory authentication](#)
- [Authentication when user mapping is being used](#)

Local authentication

Only notes common to all supported types of Windows should be referenced for Local authentication. For details, see [Notes common to all supported types of Windows on page 11-2](#).

NT domain authentication

This section contains notes that should be referenced for NT domain authentication, but does not include notes common to all supported types of Windows. For notes common to all supported types of Windows, see [Notes common to all supported types of Windows on page 11-2](#).

To use NT domain authentication without user mappings, you must register the same users as the ones already registered with the domain controller in File Services Manager, on an NIS server, or on a user authentication LDAP server. It does not matter if groups are registered under names different than those registered in the domain controller. If you register the same group with a different name, however, you must associate the group name registered in the domain controller with the group name registered in File Services Manager, on an NIS server, or on a user authentication LDAP server to refer to or set up an ACL. Therefore, we recommend that you register a group by using the same name.

If user mapping is to be used, the names of users and groups you register in File Services Manager, the NIS server, or LDAP server for user authentication, must not be the same as any user or group names registered in the domain controller. When accessing a CIFS share from an CIFS client while using a user or group name registered in the domain controller (which was set by File Services Manager, the NIS server, or LDAP server for user authentication) with an ID that differs from the user ID or group ID assigned by user mapping, folders and files might be created under the user ID or group ID set by File Services Manager, the NIS server, or LDAP server for user authentication, and not under the user ID or group ID assigned by user mapping. This occurs when the user ID or group ID has exceeded the specified range, or when the user ID or group ID has not been assigned due to a problem such as an LDAP server (for user mapping) failure.

To perform NT domain authentication without using user mappings, a CIFS client must log on to the domain to which a node belongs while accessing a shared directory created in File Services Manager. This applies even when trusted relationships exist between multiple Active Directory domains.

When user mapping is used with NT domain authentication, if authentication fails when a CIFS client attempts to access a file share, check the contents shown in [Table 5-1 Items to check if authentication fails when a CIFS client accesses a file share on page 5-7](#).

In the cases shown below, an attempt to perform NT domain authentication might fail because the node information on the domain side does not match the domain configuration information on the node side. In these cases, restart the CIFS service to recover from the state in which the CIFS client cannot connect to a CIFS share.

- A failure has occurred on the domain controller.
- You have modified the domain configuration.
- The node configuration has been modified (due to a new installation of the OS on the node, or a CIFS setting restoration when a failure occurred).

If user mappings are used, a user can access a CIFS share in an HDI system when the user belongs to a domain that has an established, trusted relationship with the domain to which the node belongs. However, if the node belongs to an Active Directory domain, a user who uses an HDI system must belong to one of the following domains:

- A domain that has a parent-child relationship with the domain to which the node belongs
- A domain that has an explicitly established one-to-one trusted relationship with the domain to which the node belongs

Active Directory authentication

This section contains notes that should be referenced for Active Directory authentication, but does not include notes common to all supported types of Windows. For notes common to all supported types of Windows, see [Notes common to all supported types of Windows on page 11-2](#).

To use Active Directory authentication without user mappings, you must have registered the same users as the ones already registered with the domain controller in File Services Manager, on an NIS server, or on a user authentication LDAP server. It does not matter if groups are registered under names different than those registered in the domain controller. If you register the same group with a different name, however, you must associate the group name registered in the domain controller with the group name registered in File Services Manager, on an NIS server, or on a user authentication LDAP server to refer to or set up an ACL. Therefore, we recommend that you register a group by using the same name.

If user mapping is to be used, the names of users and groups you register in File Services Manager, the NIS server, or LDAP server for user authentication, must not be the same as any user or group names registered in the domain controller. When accessing a CIFS share from an CIFS client while using a user or group name registered in the domain controller (which was set by File Services Manager, the NIS server, or LDAP server for user authentication) with an ID that differs from the user ID or group ID assigned by user mapping, folders and files might be created under the user ID or group ID set by File Services Manager, the NIS server, or LDAP server for user authentication, and not under the user ID or group ID assigned by user mapping. This occurs when the user ID or group ID has exceeded the specified range, or when the user ID or group ID has not been assigned due to a problem such as an LDAP server (for user mapping) failure.

To perform Active Directory authentication without using user mappings, the CIFS client must log on to the domain to which a node belongs while

accessing a shared directory created in File Services Manager. This applies even when trusted relationships exist between multiple Active Directory domains.

You can specify an Active Directory account as a user name during authentication or as a user name or group name in the File Services Manager GUI or commands. To do this, specify the name displayed for **pre-Windows 2000** in the properties of the Active Directory account. Note that the item name might differ depending on your Windows version.

With Active Directory authentication that uses user mappings, if an authentication attempt fails during a shared access from a CIFS client, see [Table 5-1 Items to check if authentication fails when a CIFS client accesses a file share on page 5-7](#).

In the cases shown below, an attempt to perform Active Directory authentication might fail because the node information on the domain side does not match the domain configuration information on the node side. In such cases, rejoin the node to the Active Directory domain to recover from the state in which the CIFS client cannot connect to a CIFS share.

- A failure has occurred on the domain controller.
- You have modified the domain configuration.
- The node configuration has been modified (due to a new installation of the OS on the node, or a CIFS setting restoration when a failure occurred).
- The HDI computer account has been modified or deleted by using the domain controller.

If you perform either of the following operations after saving system configuration information, you also need to rejoin the node to the Active Directory domain after recovering the system LU:

- Rejoin the domain by using HDI.
- Change the HDI host name.

If you modify the CIFS service authentication after saving system configuration information, you need to modify the CIFS service authentication again, save the system configuration information, and then rejoin the node to the Active Directory domain.

If an Active Directory authentication attempt fails during access from a CIFS client to a CIFS share, the authentication ticket of the CIFS client might have failed to be confirmed. Either log in again to the CIFS client machine, or restart Windows.

When you set up Active Directory authentication, make sure that the clocks of the domain controller, the HDI system, and the CIFS client are the same date and time. If there is a time discrepancy of more than five minutes among them, an authentication attempt might fail during access from the CIFS client to the HDI system.

If user mappings are used, a user can access a CIFS share in an HDI system when the user belongs to a domain that has an established, trusted

relationship with the domain to which the node belongs. However, if the node belongs to an Active Directory domain, a user who uses an HDI system must belong to one of the following domains:

- A domain that has a parent-child relationship with the domain to which the node belongs
- A domain that has an explicitly established one-to-one trusted relationship with the domain to which the node belongs

For Active Directory authentication, make sure to specify the correct domain name, user name, and password regardless of whether you are specifying a user belonging to a domain that includes the node, or a user belonging to a domain in a trusted relationship.

When an HDI node is moved to another Active Directory domain, the HDI node's computer account registered in the source Active Directory domain is sometimes not deleted and the KAQM16168-W message is output. If this occurs, authentication might fail when an attempt is made to access the HDI node from a CIFS client. Manually delete the HDI node's computer account that is still registered in the source Active Directory domain.

When Active Directory authentication is performed, the `did not have a suitable key for generating a Kerberos ticket` message might be recorded in the event log for the domain controller. This message is recorded when the encryption algorithm for Kerberos is determined, and it does not affect HDI operation. Note that the message and event ID to be recorded might differ depending on the platform of the domain controller.

Do not clear the `RC4_HMAC_MD5` check box under `Network security: Configure encryption types allowed for Kerberos in the Active Directory group policies` (Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Security Options). If this check box is cleared, the HDI system cannot join a domain and clients cannot access file shares during Active Directory authentication in HDI system version earlier than 6.4.6-00. Note that the item names might differ depending on the Windows version.

The encryption algorithms that the CIFS service supports for use with Kerberos authentication in Active Directory are `RC4-HMAC-MD5`, `AES128-CTS-HMAC-SHA1-96`, and `AES256-CTS-HMAC-SHA1-96`. The encryption algorithms that the CIFS service uses for Kerberos authentication can be set by using the `cifsopset` command. The encryption algorithms to be set must be allowed in the network security settings in `Network security: Configure encryption types allowed for Kerberos in the Active Directory group policies` (Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Security Options).

Be careful if a user who belongs to a trusted domain is subject to Active Directory authentication. In this case, the encryption algorithms allowed in `Network security: Configure encryption types allowed for Kerberos in the Active Directory group policies` (Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Security

Options) must be set as the encryption algorithms that the CIFS service uses for Kerberos authentication in Active Directory.

If Kerberos authentication in Active Directory uses the AES256-CTS-HMAC-SHA1-96 or AES128-CTS-HMAC-96 encryption algorithm to authenticate a user who belongs to a trusted domain, the trusted domain must also be set to allow use of that encryption algorithm.

If the CIFS service uses the AES128-CTS-HMAC-96 encryption algorithm for Kerberos authentication in Active Directory, depending on the Windows version, the following settings must be specified in `Network security:Configure encryption types allowed for Kerberos in the Active Directory group policies` (Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Security Options):

- Allow RC4-HMAC-MD5 and AES128-CTS-HMAC-96.
- Do not allow AES256-CTS-HMAC-96.

If you change the encryption algorithm used for Kerberos authentication in Active Directory and the CIFS service is restarted, the CIFS client must log off and then log in again.

This note applies if you add encryption algorithms allowed in `Network security:Configure encryption types allowed for Kerberos in the Active Directory group policies` (Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\Security Options), and set the CIFS service to use those encryption algorithms for Kerberos authentication in Active Directory. In this case, again set the password of the domain user who is used to join an Active Directory domain, and then have the HDI system rejoin the Active Directory domain.

Make sure that the domain controller supports the versions of the SMB protocol that the CIFS service uses to communicate with the domain controller. If it does not support those versions of the SMB protocol, the HDI system cannot join a domain and clients cannot access file shares during Active Directory authentication. In this case, change the values specified for the options `client_ipc_max_protocol` and `client_ipc_min_protocol`, of the command `cifsoptset`, that determine the versions of the SMB protocol that the CIFS service uses to communicate with the domain controller.

If Active Directory authentication is used in HDI system version earlier than 6.4.2-00, HDI system will establish an anonymous connection to the domain controller after a failover is performed. For this reason, if anonymous connections are prohibited by the domain controller settings, HDI system will be unable to access the domain controller after a failover. If this problem occurs, set the domain controller to permit anonymous connections.

If you change the UNIX attributes of users and groups on the domain controller, manually replace the ACLs for files and folders in HDI system and, on the **CIFS Service Maintenance** page of the **Access Protocol Configuration** dialog box in File Services Manager, click the **Clear User Map Cache File** button to delete the user mapping cache file.

In a cluster configuration, use the same user mapping settings for both nodes.

In the OUs (organizational units), set the access permissions for the HDI system computer accounts and for the domain user accounts specified in HDI system.

If you are operating a WORM file system and if the UIDs or GIDs of a domain's users change (for example, because of an Active Directory migration), users might become unable to access WORM files, depending on the ACL settings. It is also impossible to match the pre-migration UIDs and GIDs to the post-migration UIDs and GIDs. Therefore, if you are using RID user mapping, set the access permissions for files to `Everyone` before converting the files into WORM files. If you use an Active Directory schema for mapping, however, Active Directory settings can be used to match UIDs and GIDs to the IDs used before the migration. This also applies to LDAP user mapping.

If you set the Read Only Domain Controller (RODC) as an authentication server, Kerberos tickets will not be issued and you might be unable to access CIFS.

While HDI system is being replaced, if two or more HDI nodes have the same host name, do not let those nodes participate in Active Directory under the same host name.

Do not block communications between HDI system and Active Directory (including trusted entities) by using firewalls or similar devices.

While NTLM authentication is in use, do not block NTLM authentication in Active Directory.

Specify LDAP signature settings for Active Directory that match the settings for HDI system.

Authentication when user mapping is being used

When user mapping is used, if authentication fails in accessing a file share from a CIFS client, check the contents shown in the following table.

Table 5-1 Items to check if authentication fails when a CIFS client accesses a file share

#	Item to check	Contents to check	Action
1	Range of UIDs and Range of GIDs , shown in the CIFS Service Maintenance page of the Access Protocol Configuration dialog box*1 *4 *6	If the entire range of user IDs or group IDs is in use.	If so, expand the range of the user IDs or the group IDs.
		Part of the range of the user IDs or group IDs is not in use.	Check item #2.
2	The operating status of the CIFS service, shown in the List of Services page of the Access	The CIFS service is running properly.	Check item #3.

#	Item to check	Contents to check	Action
	Protocol Configuration dialog box*2	The CIFS service is not running properly.	Restart the service.
3	The LDAP server set in the CIFS service configuration definition*4	The LDAP server is running properly.	Depends on the LDAP operational status.
		The user ID and group ID are not registered for a user whose authentication has failed.*7	Register the user ID and group ID.
4	View the CIFS log*3 (/var/log/cifs/log.winbindd)	Whether failure information is output.	Act as indicated in the output log.
5	Execution of the <code>umapidget</code> command*5	The user ID and group ID of the user that could not be authenticated is outside the valid range.	Expand the range of the user IDs or the group IDs.
		The user ID and group ID of the user that could not be authenticated is in the valid range.	Check item #2 in this table.
6	The Active Directory management window for the user whose authentication has failed*8	The user ID and group ID on the UNIX Attributes tab are not registered for the user whose authentication has failed.	Register the user ID and group ID on the UNIX Attributes tab.

*1:

Check the allocation status of the user-mapped user IDs and group IDs.

*2:

Confirm that a failure did not occur in referencing or allocating user-mapped user IDs or user-mapped group IDs. For details on how to check the operating status of the CIFS service, see the *Administrator's Guide*.

*3:

For details about how to view log files (for example, the CIFS log), see the *Administrator's Guide*.

*4:

This item applies only when **Use user mapping using LDAP.** is selected as the user mapping method.

*5:

This item applies only when **Use user mapping using RIDs.** is selected as the user mapping method.

*6:

This applies only when you select **Allocate automatically** as an ID allocation method.

*7:

This applies only when you select **Allocate manually** as an ID allocation method.

*8:

This applies only when you select **Use user mapping using Active Directory schema** as a user mapping method.

Whenever a connection to the LDAP server used for user mapping fails, the LDAP server cannot be accessed again for at least five minutes. Users (such as a new domain user or a domain user who has deleted cache files) who usually access the LDAP server by using the CIFS service will not be able to access the LDAP server. Correct the problem that is preventing connections to the LDAP server, either wait five minutes or restart the CIFS service, and then attempt to access the CIFS service again. If the LDAP server is restarted after the problem is corrected, restart the LDAP server, and then restart the CIFS service.

When an error occurs on a network with a domain controller, and when you receive CIFS-service-related error information using SNMP or email notification from a node, the node cannot acquire user and group information from the domain controller for five minutes after failure detection. Accordingly, the node fails to authenticate users during this five-minute period. If this problem occurs, correct the error that is preventing a connection to the domain controller, and then either wait five minutes or restart the CIFS service. After that, access the CIFS service.

If multiple CIFS clients try to connect to the HDI system in a short period of time one or more of, the CIFS clients might fail to connect to the HDI system, depending on the current load on the HDI system, the processing power of the CIFS clients and the DC server, and the network environment. In such cases, we recommend that you do the following:

- Make a connection prior to accessing HDI with CIFS.
If the CIFS client makes a connection prior to connecting to HDI, you can better spread out the load on the DC server or the network. To make a connection prior to connecting to HDI, use the Windows API `WNetAddConnection2()` function or the `net` command. (If you use the `net` command to make a connection beforehand, authentication might be necessary again when the CIFS access is attempted.) Also, to avoid being disconnected from an HDI timeout, set the HDI timeout value to 0. (0 indicates no timeout.) In File Services Manager, you can set the timeout value in **Client time-out** on the **CIFS Service Management** page (**Setting Type:** *Performance*) in the **Access Protocol Configuration** dialog box.
- If you fail to connect to HDI, try to make a connection on the CIFS client again. When attempting to reconnect to HDI, we recommend that you

wait approximately 30 to 60 seconds, in order to better spread out the load on the DC server and the network.

Reference: In the domain environment where HDI resides, when the DC server and the network are operating normally, the number of connections that can be processed per second is approximately 100 for NTLM authentication and approximately 10 to 12 for Kerberos authentication. (Kerberos authentication takes more time to process authentication because it prevents reply attacks.)

Procedure for Migrating User Resources in a Windows Domain Environment

This chapter contains notes on the migration of user resources created in the Windows domain environment and also how to migrate user resources onto the HDI system by using the backup utility.

- [Before performing resource migration](#)
- [Using the backup utility to perform migration](#)

Before performing resource migration

An access control list (ACL) provided for a CIFS share on the HDI system differs depending on the type of file system (Classic ACL type or Advanced ACL type). A Classic ACL type file system maps a UNIX ACL compliant with POSIX to a Windows ACL. UNIX ACL, mapped to a Windows ACL. A UNIX ACL basically resembles a Windows ACL, but since it is based on UNIX file permissions, there are portions with significant functional differences. Therefore, some of the Windows ACL functions are not available. With an Advanced ACL type file system, you can set up advanced access permissions similar to a Windows ACL, which allows for more Windows-like access control. For details about ACLs, see [Access Control Lists on page 8-5](#) in this manual.

Migrating from a Windows file server to an HDI system

The following table lists the effects of file system differences when a file system is migrated from a Windows environment to an HDI system.

Table 6-1 Effects of file system differences when a file system is migrated from Windows to HDI

#	(In Windows) Items		(In HDI) Classic ACL type	(In HDI) Advanced ACL type	Remarks
1	Operator		CIFS administrator registered in File Services Manager	CIFS administrator registered in File Services Manager	For a non-administrative user, the ACL of the folder to be migrated to, files to be migrated, and the ACL of the folder might affect the migration of data, change of ownership, and the ACL configuration.
2	Migration of the owner	User	Migratable	Migratable	The SID and UID must be resolved.#1
		Group	Not migratable	Migratable	The SID and GID must be resolved.#1
		Built-in/Well-known SID account	Not migratable	Not migratable	--
		SID unresolvable account	Not migratable	Not migratable	The SID and GID must be resolved.#1
3	Migration of ACLs#2 #3	File owner	Migratable	Migratable	--
		Primary group	Migratable	Migratable	--

#	(In Windows) Items	(In HDI) Classic ACL type	(In HDI) Advanced ACL type	Remarks
		DACL (Discretionary ACL)	Partly not migratable ^{#4} ^{#5} (mapped to POSIX ACL)	Migratable The SID, UID, and GID must be resolved. ^{#1}
		SACL (System ACL)	Not migratable ^{#6}	Not migratable ^{#6} --
4	Migration of file attributes	Read-only attribute	Migratable	Migratable --
		Archive attribute	Migratable ^{#6}	Migratable The Windows backup utility can migrate attributes from the migration source, but the XCOPY command adds the archive attribute to all files.
		System attributes	Migratable ^{#6}	Migratable --
		Hidden attribute	Migratable ^{#6}	Migratable --
		Directory attribute	Migratable	Migratable --
		Encryption attribute	Not migratable ^{#6}	Not migratable ^{#6} The XCOPY command decrypts files. The Windows backup utility causes an error if an attempt is made to restore an encrypted file.
		Compress attribute	Not migratable ^{#6}	Not migratable ^{#6} The compress file attribute is disabled, causing uncompressed data to be saved.
		Offline attribute	Not migratable	Not migratable --
		Normal file attribute	Migratable	Migratable --
		Temporary file attribute	Not migratable	Not migratable --
	Sparse file attribute	Not migratable	Not migratable --	

#	(In Windows) Items	(In HDI) Classic ACL type	(In HDI) Advanced ACL type	Remarks	
	Reparse point attribute	Not migratable	Not migratable	--	
	Non-indexed file attribute	Not migratable	Not migratable	--	
5	Migration of timestamps	Access time	Not migratable	Not migratable	The date and time of the copy (migration) operation is set.
		Modified time	Migratable	Migratable	--
		Created time	Not migratable ^{#7}	Not migratable ^{#7}	--

Legend: --: No remarks

#1:

For example, if Windows Active Directory Migration Tool (ADMT) is used to migrate user accounts in an Active Directory domain, file and folder access privileges might not migrate from the source file server to HDI because the (SID) accounts used in the source file server cannot be resolved by HDI. To migrate file and folder access privileges, reassign the original account information to the account information in the migration destination before starting the migration. Note that HDI provides the steps for mapping the original account information when data is imported from other file servers. For details about how to import data from other file servers by using the CIFS protocol, see *Administrator's Guide*.

In addition, if you used the `cifsoptset` command to enable ACL for the default Windows-domain user groups `Authenticated Users` and `Network`, `Authenticated Users` and `Network` can be recognized. The ACL settings for `Authenticated Users` and `Network` are not imported at this time, even if you import data from another file server. In such cases, use the mapping function of the migration-source account, in the same way as for other accounts. For details on how to enable or disable ACL for the default Windows-domain user groups `Authenticated Users` and `Network`, see the *CLI Administrator's Guide*.

#2:

The ACL revision that can be specified is `ACL_REVISION (0x2)`. The ACL revision of `ACL_REVISION_DS (0x4)` can not be specified.

#3:

The ACE types that you can specify are only `0x0 (Access allowed)` and `0x1 (Access denied)`. Do not specify other ACE types.

#4:

In the Classic ACL type file system, migration might not be possible if the number of ACEs exceeds 63, including the owner and the group. In the Advanced ACL type file system, migration might not be possible if the number of ACEs exceeds 700.

#5:

If a discretionary access control list (DACL) includes the following items, the access control entries (ACEs) of such items are not migrated:

- Users or groups that cannot be recognized in the HDI system
- BUILTIN/well-known SID accounts other than Everyone, CREATOR OWNER, and CREATOR GROUP.

#6:

HDI does not support this. Because HDI systems do not support SACLs, HDI systems cannot perform the same audit functionality as the standard audit functionality in Windows. To perform auditing, consider using the CIFS access logs instead.

#7:

Migration is possible only when the file system to be migrated to has been configured to record the created time of files. When the file system is not configured to record the created time of files, the created time is either the modified time, the access time, or the attribute modified time of the file, whichever is the oldest.

The following table shows precautions for using Windows-standard commands or programs to migrate user resources. Considering the backup utility's ability to migrate user resource attributes, we recommend that you use this utility on an HDI system.

Table 6-2 Using commands or applications to migrate user resources

#	Command / application	Precautions regarding migration
1	Copying via Explorer	ACL information cannot be recovered. ACL information contains file owner and primary group information, and the resource owner is the user that performed file migration. Therefore, the ACL needs to be set again after file migration.
2	XCOPY	By using the XCOPY command with certain options, a CIFS administrator registered in File Services Manager can migrate owner and ACL information. If the owner is not a user who can be recognized in the HDI system ^{#1} , the user cannot migrate the file.
3	Backup utility (Windows standard) ^{#2}	By using the backup utility, a CIFS administrator registered in File Services Manager can migrate owner and ACL information. Information about the files that have been backed up is output to a report.

#1:

A user who can be recognized in the HDI system means a user who has a user name to which SID can be mapped on the HDI system. Therefore, users and groups registered in the domain to which they belong can be recognized in the HDI system. Users and groups specific to Windows clients (including built-in users) except `Everyone`, `CREATOR OWNER`, and `CREATOR GROUP` cannot be recognized in the HDI system. However, if users and groups of the domain have already been deleted from the domain prior to migration, the HDI system cannot recognize those users and groups.

Supplementary note:

If you used the `cifsoptset` command to enable ACL for the default Windows-domain user groups `Authenticated Users` and `Network`, `Authenticated Users` and `Network` can be recognized. For details on how to enable or disable ACL for the default Windows-domain user groups `Authenticated Users` and `Network`, see the *CLI Administrator's Guide*.

#2:

The backup utility here means the Windows standard backup tool.

Notes on migrating files by using the `XCOPY` command or the backup utility:

- To migrate an ACL associated with a user resource, the Windows host to be migrated from, the Windows host on which the command or application runs, and the HDI node to be migrated to must belong to the same Windows domain, and at the same time, the CIFS service configuration definition on the HDI system must use user mappings.
- Only a CIFS administrator registered in File Services Manager can migrate backed-up files to the HDI system. For details on how to register a CIFS administrator, see the *Administrator's Guide*.
- If the authentication method for the CIFS service configuration definition is anything other than NT domain authentication or Active Directory authentication, or is NT domain authentication or Active Directory authentication but user mapping is not used, ACLs set in user resources cannot be migrated. When migration is performed, the owner of the migrated user resources is `root`, and the group is the group to which the user who performed user resource migration belongs.
- For details on whether you can migrate ACLs (on the Windows host), file attributes, ACLs, and timestamps (on the HDI system), see [Table 6-1 Effects of file system differences when a file system is migrated from Windows to HDI on page 6-2](#).

Notes on migrating a file or folder that has 64 or more ACEs (Classic ACL type):

When a file or folder has 64 or more ACEs, including the owner and group of the file or folder, you might be unable to migrate all the ACEs from the Windows server to the HDI system. In this case, you can avoid this problem as follows: Specify that the users with the same ACL belong to

the same group, and then set the ACL to the applicable group. By doing this, you can reduce the number of ACEs to 63 or fewer.

Notes on migrating a file or folder that has 701 or more ACEs (Advanced ACL type):

When a file or folder has 701 or more ACEs, including the owner and group of the file or folder, you might be unable to migrate all the ACEs from the Windows server, to the HDI system. In this case, you can avoid this problem as follows: If there is more than one user to whom the same access permissions are set in the ACL, make those users belong to a single group, and set the ACL to this group, not to the multiple users, to reduce the number of ACEs to 700 or fewer.

Notes on ACEs to be added during migration (Classic ACL type):

During the migration of user resources, the following ACEs are automatically added to each file and folder:

- File migration
The owner and group of the applicable file is set to and indicated in the ACL.
- Folder migration
When a folder has an ACL to be applied to its subfolders and files, ACLs `CREATOR OWNER` and `CREATOR GROUP` are additionally set and indicated.

Notes on file migration when the file owner is not a domain user that the HDI system can recognize:

When you use the backup utility, the owner of the file to be migrated to is either of the following:

- CIFS administrator (root user)
- User included in the ACL of the file to be migrated from if the ACL of the file contains a combination of the user and the primary group

When you use the `XCOPY` command, the migration of the file will fail with an error, creating an empty file at the intended destination of the migration.

Notes on file migration when the ACL of a user other than the file owner contains a user or group with an SID that the HDI system cannot recognize:

During migration, if there is an ACE of a user or group with an SID that the HDI system cannot recognize, the ACL is migrated without the ACE.

Using the backup utility to perform migration

The following is an overview of migrating user resources from the Windows domain environment to an HDI system.

1. Obtaining the file attributes and ACL information of the file to be backed up

When you migrate from the Windows domain environment to the HDI system, you might have to reconfigure the file attributes and the ACL

after the migration, due to the differences in file attributes and ACL specifications. Therefore, use the `CACLS` command and the `ATTRIB` command before migration to obtain the file attributes and ACL information for the file to be backed up.

2. Creating a backup file

Use the backup utility to create a file that contains a backup of the data you want to migrate. For details about how to create a backup, see Help or other documentation for the backup utility.

When the backup file has been created, confirm that the folders and files to be migrated have been placed correctly in the file.

3. Registering a CIFS administrator into File Services Manager

You must migrate backed-up files to the HDI system as a CIFS administrator. If a user other than CIFS administrator (file owner) migrates it, the file attribute may not be migrated properly.

You must register either a user who performs file migration from the Windows server or register a group to which the user belongs as a CIFS administrator in File Services Manager. You can do this in the **CIFS Service Management** page (**Setting Type:** Administration) of the **Access Protocol Configuration** dialog box. For details, see the *Administrator's Guide*.

4. Creating a file system and a CIFS share

Create the file system and a CIFS share to be migrated to on the HDI system.

To do this, in File Services Manager, either create the file system and a CIFS share at the same time in the **Create and Share File System** dialog box, or first create the file system in the **Create File System** dialog box and then add a CIFS share in the **Add Share** dialog box. For details, see the *Administrator's Guide*.

Note

HDI CIFS service configuration definitions are not case sensitive by default. This is done to unify operations with those of Windows servers. By executing the `cifsoplist` or `cifsopset` command from an HDI node, configuration settings for the CIFS service or individual CIFS shares in the node can be viewed and modified.

Migration performance might be improved if, before performing a migration to an HDI system (a backup file recovery operation), you temporarily change the setting regarding case sensitivity with file names. However, if you do so, after the migration is finished, be sure to return the setting to its default before performing operations in an HDI system. If you perform operations in an HDI system while the setting is set to be case sensitive with file names, and there are multiple files within an HDI shared directory that, other than case, have the same file name, you might accidentally use a file you did not intend to use because the CIFS client is not differentiating between upper case and lower case.

5. Restoring the backup file

Restore the data in the backup file that was created by using the backup utility in step 2 to the file share created in step 4.

The CIFS administrator registered in step 3 must log on to the Windows system that contains the backup file, and then restore the data. For details about the restore operation, see Help or other documentation for the backup utility.

6. Checking the CIFS logs

If the user mapping functionality does not work correctly during file migration, the file owners and ACL settings might not be migrated correctly. After file migration is completed, use the CIFS logs (`/var/log/cifs/log.winbindd`) to check whether an error has occurred in the user mapping functionality. If an error has occurred, correct it and then perform the migration again.

For details about how to view log files (for example, the CIFS log), see the *Administrator's Guide*. For details on messages in the CIFS log (`/var/log/cifs/log.winbindd`), see [log.winbindd on page A-5](#).

7. Reconfiguring the ACL settings of the migrated file

As we have seen before, when you migrate from the Windows domain environment to the HDI system, you might be unable to restore ACLs correctly due to differences in file attributes and ACL specifications. In this case, reconfigure the ACLs based on the ACL specifications on the HDI system.

Only a CIFS administrator registered in File Services Manager can reconfigure ACLs. As the CIFS administrator registered in step 3, access the CIFS share and reconfigure the ACL settings.

Accessing CIFS Shares

This chapter describes the procedure for accessing a shared directory from a CIFS client, and also provides related notes.

- [Access method](#)
- [Notes on access from CIFS client](#)
- [Notes on CIFS access in an environment where Anti-Virus Enabler is applied](#)
- [Setting home drives](#)
- [Notes on using the Windows roaming user profile functionality](#)

Access method

To allow a CIFS client to access a shared directory, specify either of the following paths. Note that the same specification format must be used for each user or for each system.

- Path 1: `\\physical-node-host-name#\CIFS-share-name\path-to-a-directory-to-be-used`
- Path 2: `\\virtual-IP-address\CIFS-share-name\path-to-a-directory-to-be-used`

#

The physical node host name in the specified path is equivalent to either the host name or NetBIOS name of an HDI node. You cannot specify, for the *physical-node-host-name*, the alias registered in the CNAME record in the DNS.

To access a CIFS share via an IPv6 connection, you must specify a host name or an `ipv6-literal.net` name for the physical node host name or virtual IP address. An `ipv6-literal.net` name is an IP address in a format in which the delimiting colons (:) of an IPv6 address are replaced with hyphens (-), and `.ipv6-literal.net` is added to the end of the address, as follows:

```
ipv6-literal.net name when the IPv6 address is fd00::5:50
    fd00--5-50.ipv6-literal.net
```

Client connection method

When NetBIOS over TCP/IP is enabled as the client connection method, Windows tries both NetBIOS over TCP/IP (port 139) and Direct Hosting of SMB (port 445) at the same time (in parallel) for the CIFS connection, and uses the connection that was established first.

This behavior is also explained on the following Microsoft Web page:

<http://support.microsoft.com/kb/204279/en-us>

The connection not selected is immediately closed by the client. However, depending on the timing of the disconnection, since a child process of `smbd` is already attempting to respond to the request, a message indicating disconnection from the client might be logged depending on the disconnection. (For details on the meanings of messages, see [CIFS logs on page A-3](#) and [log.smbd on page A-3](#).) This message does not affect the CIFS access itself because communication with the client occurs on the connection established first.

Name resolution services

The CIFS client can use WINS, DNS, `lmhosts`, and other name resolution services. The following table lists notes on using these name resolution services.

Table 7-1 Notes on using name resolution services

Name resolution service	Note
WINS	Define all the other CIFS clients on the network as WINS clients. Manually register the virtual IP address, and the host name or NetBIOS name of an HDI node to the WINS server.
DNS	Manually register the virtual IP address, and the host name or NetBIOS name of an HDI node to the DNS server.
lmhosts	Register the virtual IP address, and the host name or a NetBIOS name of an HDI node to <code>lmhosts</code> of all the CIFS clients.

Notes on access from CIFS client

This subsection provides notes on access from CIFS clients.

- Notes on connections are as follows:
The maximum number of CIFS clients that can connect to the CIFS service is 12,000 (see [Table 7-2 Maximum number of CIFS client connections and the maximum number of CIFS shares on page 7-4](#) for details). This maximum does not change even when a failover has forced multiple resource groups to operate on one node. If the maximum number of CIFS clients that can be connected to the CIFS service has already been reached when a CIFS client attempts to connect to the CIFS service, an error message is displayed on the client.

A connection to the CIFS service is not immediately closed when access to the CIFS share stops. The connection can be closed by using either of the following methods:

- Logging in to the CIFS client again
- Disconnect all the connections to the CIFS share on the HDI system.

In addition, when a file is not open, the HDI system automatically disconnects the connection to the CIFS client if a CIFS client does not access the HDI system for a set period of time, specified in **Client time-out**. For details on the time specified in **Client time-out**, see the *Administrator's Guide*.

If a user tries to access a CIFS share that has been disconnected by the HDI system, the CIFS client automatically tries to reconnect to the CIFS service. Accordingly, the user does not need to enter authentication information again for reconnection. However, if the maximum number of CIFS clients that can be connected to the CIFS service has already been reached, as it does with an ordinary connection, the attempt to access the CIFS share fails.

When you display the content of a CIFS share in the Explorer menu, Explorer regularly accesses the CIFS service. Therefore, the HDI system does not automatically disconnect the connection.

Table 7-2 Maximum number of CIFS client connections and the maximum number of CIFS shares

Node model	Automatic reloading	Maximum number of CIFS client connections ^{#1}	Maximum number of CIFS shares ^{#1}
D51B-2U	N	300	7,500
Single-node configuration (node not connected to a storage system)	Y	300	256
D51B-2U	N	12,000	7,500
Cluster configuration (node connected to a storage system)	Y	4,800	256
Compute Rack 220 or Compute Rack 220S	N	300	7,500
Single-node configuration (node not connected to a storage system)	Y	300	256
Compute Rack 220	N	12,000	7,500
Cluster configuration (node connected to a storage system)	Y	4,800	256
Compute Rack 210H	N	6,000	7,500
Cluster configuration (node connected to a storage system)	Y	2,000	256
HDI Remote Server Centrally Managed configuration ^{#2}	Y ^{#3}	30	256
HDI Remote Server	N	30	256
Locally Managed configuration ^{#4}	Y	45	256

Legend:

Y: Automatic reloading is enabled.

N: Automatic reloading is disabled.

#1

The maximum number of CIFS client connections and CIFS shares are the number per cluster in a cluster configuration, or the number per node in a single-node configuration.

#2

In this configuration, a single node of the HDI system is linked to an HCP system or HCP Anywhere.

#3

For a centrally managed configuration of HDI Remote Server, automatic reloading is always enabled because it is the default setting of HCP Anywhere.

#4

In this configuration, a single node of the HDI system is linked to an HCP system, and is not linked to HCP Anywhere.

- If a CIFS client has cached a write request, and then either a failure occurs or the disk space becomes insufficient on either the CIFS client or the network, the cache might be unable to guarantee the data. (For example, the writing of a file seems to have succeeded, but the data was not properly written.) Keep this in mind if you set up the CIFS client to cache file updates for a CIFS file share.
- If the setting is enabled so that file update data is cached on a client, and write requests are being cached by a specific CIFS client, before another CIFS client can successfully open the same file, the processing to flush the cache data from the first client must first be performed. In addition, if the setting for writing data synchronously for write requests and close requests is enabled, and a different CIFS client attempts to open the same file, the other CIFS client will be unable to open the file until write processing and the processing to flush the data to a disk drive from the first client are complete. Therefore, the other CIFS client might need to wait for some time until the file can be opened.
- When an update is performed by multiple clients for the same file in a CIFS share, if the updated data for the file is set to be cached on the client, an access delay or deterioration in data reliability might occur. Therefore, we recommend that you save files that might be accessed by multiple clients in a CIFS share that is set to not cache on the client.
- If an operation is performed from a CIFS client to move a folder within a CIFS file share, the update date and time for the moved folder are changed to the move date and time.
- When a client creates files or folders in a CIFS share, the larger the number of files and folders to be stored in the same folder at the destination, the longer the time required for creating them. This is because the system determines the case of the characters in the names of the file and folder to be created, for checking for any duplications.

In the HDI system, the initial setting of the CIFS service configuration definition is `insensitive` to the case of the characters in the file names, in order to follow the behavior of the Windows server. For example, `ABC.txt` and `abc.txt` are determined to be the same and cannot be created in the same folder.

If the number of files and folders in the same folder exceeds 1,000, the duplication check will take far more time. Make sure that a single folder does not contain an excessively large number of files and folders. As the number of files and folders in the same folder increases, a greater load is placed on the HDI system, and displaying files and folders takes more time. Therefore, make sure that a single folder does not contain an excessive number of files and folders.

Reduction in the processing time can be expected if the `cifsoptset` command is used to change the setting of the CIFS service configuration definition to `sensitive` to the case of the characters in the file names. Before doing so, make sure that there are no file names in the same folder that will become duplicated without case sensitivity. If there

are any duplications, a CIFS client might perform operations on an unintended file.

- If you use SMB signing for communication with a CIFS client, you can prevent man-in-the-middle attacks that tamper with SMB packets being transferred. Note, however, that the security improvements granted by SMB signing will also degrade file access performance. Before you can use SMB signing, the necessary settings must be specified for both the client and the HDI system. The HDI system always uses SMB signing when the client requests SMB signing for communication via the SMB 2.0, SMB 2.1, or SMB 3.0 protocol. In addition, you can use the `cifsoptset` command to specify whether to use SMB signing for SMB 1.0 communication. With the initial settings, the HDI system does not use SMB signing for SMB 1.0 communication.
- If the system administrator is changing the CIFS service configuration definition while you are performing an operation on a file system from a CIFS client, that operation might fail. If the operation fails, wait until the CIFS service configuration definition change has been completed, and then retry the operation.
- If the system administrator changes the information on a CIFS file share in a file system while a CIFS client is using the file system, the changes might not be applied. The CIFS client user must reconnect to the CIFS file share or restart Windows to apply the changes.
- If a failover occurs, an operation performed from a CIFS client that is using the services of the resource group moved by the failover or a failback is forcibly canceled.
- If access and write operations from CIFS clients are suppressed on a node that is accessed by a CIFS client, and if the CIFS service is subsequently restarted, the CIFS service is placed in the incomplete status. In this case, the **List of Services** page of the **Access Protocol Configuration** dialog box might display **Running in Status** and **The service is incomplete. Restart the service.** in **Information**, and the connection to the CIFS share might be disabled.
- When multiple users share the same file, contention between Windows applications can occur, causing the ACL set individually for a file to be lost due to the Windows application specifications.
- When an error has occurred in a file system, you might not be able to view any of the files and directories in a CIFS share in the file system. However, when the error has been corrected, the contents of the CIFS share can be viewed as usual.
- When a CIFS client accesses a stub file, it might take some time to process the file. Also, a timeout might occur. If **Parallel** is specified as the CIFS service configuration definition or the CIFS share attribute in **Windows(R) client access policy** and a CIFS client accesses a stub file, the timeout period is set to a maximum of 15 minutes. For details on stub files, see the *Installation and Configuration Guide*.
- When an attempt is made from a CIFS client to delete a stub file in a file system that is linked to a Hitachi Content Platform (HCP) system, as seen from the CIFS client, the file might seem to have been deleted, even

though the file has not actually been deleted, because an HCP error occurred. Therefore, if you want to delete a stub file in a file system that is linked to an HCP system, after attempting to delete the file, make sure that the file has actually been deleted by checking the folder the file is stored in.

- If a file system is re-created after recovering from an LU failure, data that has been migrated to an HCP system can be restored to an HDI file system. For details on how to restore a file system whose data has been migrated to an HCP system, see the *Troubleshooting Guide*.

Due to the amount of time it takes to completely restore a file system, there is the possibility that a client might unknowingly access data that has not yet been restored. This can cause access attempts to fail due to timeouts occurring from the CIFS client. Timeouts occur when it takes too long to display the intended files because the parent directory contains a large amount of data. If a network error or some other error is displayed on the CIFS client, wait a while, and then try to access the file again.

- When the OS on the node is under a heavy load, a CIFS client attempting to access a CIFS share might generate an error indicating that no space is left on the disk even before the file system capacity reaches 100%.
- Files and folders used by the system (see the table below) might be displayed on CIFS clients. The following table provides notes on the files and folders used by the system.

Table 7-3 Notes on the files and folders used by the system

File name or folder name	Note
.arc	This folder is created if one of the following occurs: <ul style="list-style-type: none"> • HDI data is migrated to an HCP system. • Data is imported from another file server. Do not edit or delete data under this folder. If data under this folder is edited or deleted, system information might become inconsistent.
.backupdates	This file is created when the NDMP functionality is used. This file cannot be edited or deleted.
.conflict	This folder is created if data contention occurs between HDI systems that share data on an HCP system. Data under this folder cannot be edited or deleted.
.conflict_longpath	When files in conflict in HDI systems that share data on an HCP system are set to be saved in the directory in which the files in conflict were originally stored, this folder is created if a file conflict occurs for a file whose path length exceeds 187 characters. Data under this folder cannot be edited or deleted.
.history	This folder is created when the file version restore functionality is used. Data under this folder cannot be edited or deleted.
lost+found	This folder is created when the integrity of the file system is checked. Data under this folder cannot be edited or deleted.

File name or folder name	Note
<code>.lost+found</code>	This folder is created if there is an inconsistency between files in an HDI system and files in a migrated HCP system. Data under this folder cannot be edited or deleted.
<code>.system_reorganize</code>	This folder is used for temporarily saving data to allocate the unused capacity of the inode area. This folder is always created at system installation. Data under this folder cannot be edited or deleted.
<code>.temp_backupdates</code>	This file is created when the NDMP functionality is used. This file cannot be edited or deleted.

- During batch configuration of ACLs for objects in a folder, if the objects include a file or folder named `.snaps` or a file or folder described in [Table 7-3 Notes on the files and folders used by the system on page 7-7](#), the ACL setting fails and processing is interrupted. You can prevent this by using the `cifsoptset` command to specify a setting so that files and folders used by the system are excluded from the list. After finishing the ACL setting, change the setting again so that files and folders used by the system are included in the list.

Alternatively, before setting the ACLs, you can create a CIFS share for the folder under a different name and execute the `cifsoptset` command on the share to specify a setting so that files and folders used by the system are excluded from the list.

For details about how to set whether files and folders used by the system are eliminated from the list, see the *CLI Administrator's Guide*.
- If you specify a setting so that files and folders used by the system are excluded from the list or if you enable Access Based Enumeration (ABE), an attempt to delete the folder that stores a file or folder that is not displayed fails. At this time, no error is displayed and the folder appears to be deleted from a CIFS client. However, if you refresh the window, the target folder is displayed.
- Generations of the data migrated to an HCP system are managed for each date and time when migration was performed by using version management (versioning). In HDI systems, you can use the data whose generations are managed to re-create the directory structure at the time when a migration was performed. By opening a directory that was re-created in the `.history` directory under the shared directory to HDI clients, you can restore data in file units, even if a user accidentally deletes a file, or if a user who shares data with other HDI systems via the linked HCP system overwrites an updated file (File version restore functionality). For details about the file version restore functionality, see the *Installation and Configuration Guide*.
- In a home-directory-roaming file system, if a user updates the files when the HDI systems cannot be synchronized due to a failure such as a communication error, a directory named `.conflict` might be created in the home directory of the user and then the updated files might be saved in the directory.

In a read-write-content-sharing file system, if the same file is updated on multiple HDI systems before the HDI systems are synchronized, a conflict will occur the next time the HDI systems are synchronized. In an HDI system where a conflict occurred, files are saved in the directory in which the files in conflict were originally stored or in the `.conflict` directory just below the mount point of the read-write-content-sharing file system. The end user must check whether the necessary files are saved in the directory that the system administrator told them about. If the files in conflict are not stored in the directory, you must contact the system administrator.

To use the files stored in the `.conflict` directory, copy each file individually to a location of your choosing other than the `.conflict` directory. If the files are copied by directory, incorrect access permissions might be set.

For end users to access the `.conflict` directory, it is necessary to configure settings on the client side so that all files and folders are displayed.

- In the read-write-content-sharing file system, an I/O error might occur if a directory is manipulated remotely from a different site or if a failover, failback, or other temporary failure occurs when a directory is manipulated. If an I/O error occurs, wait for a while, and then try again. If the error recurs, contact the HDI administrator.
- If the file owner uses Explorer to change the access permissions of a file that cannot be accessed, an error might occur, and the property window might become inoperable. For this reason, do not change access permissions of files for which you do not have the necessary access permissions. If you cannot access a file, the access permissions must be changed by a CIFS administrator or a user who has the necessary permissions to access the file.
- When files or folders are migrated by using the `robocopy` command, the SACL (audit ACL) is not migrated regardless of whether an error occurs in the `robocopy` command.
- If you are using the Windows API to access CIFS share files, when opening files from a client by using `CreateFile()`, specify an access mask that contains `FILE_READ_DATA`, and then open the file. If you specify `FILE_WRITE_DATA` for the access mask, open a file for which `FILE_READ_DATA` is not specified, and then lock the file without specifying `LOCKFILE_EXCLUSIVE_LOCK` for the file handle by using `LockFileEx()`, `LockFileEx()` results in an error (`ERROR_INVALID_HANDLE`). The same behavior applies when you specify any of the following: `GENERIC_READ`, and `GENERIC_WRITE`.
- Access from a multifunction client, such as one with a printer, is not supported.
- Privileges such as the privilege "Restore files and directories" can be granted to Windows user accounts; however, access to files and folders on which such privileges are used cannot be restricted. For this reason, avoid using such privileges. To access all files and folders, have a CIFS administrator perform the relevant operations.

Notes on CIFS access in an environment where Anti-Virus Enabler is applied

If a file being accessed in a CIFS share has been infected with a virus or if an error occurs during real-time virus scanning, the result of the operation might not be the intended result. For example, if a file to be stored in a CIFS share is infected with a virus, the file will not be stored.

In an environment where Anti-Virus Enabler is used, depending on the operating conditions of the CIFS client, sessions might be disconnected due to a CIFS client timeout, and application programs might end abnormally. If this error occurs, the following messages are output to the CIFS client.

Messages

```
The error conditions of the CIFS client in an Anti-Virus Enabler
environment are as follows:
Err_no=6 err_msg=The handle is invalid.
Err_no= 64 err_msg=The specified network name is no longer available.
Err_no=121 err_msg=The semaphore timeout period has expired.
```

If a CIFS client request causes the above errors in an Anti-Virus Enabler environment, the CIFS log that has detected the session disconnection, due to the client timeout, might indicate the following error conditions:

Example in `/var/log/samba/log.smbd`:

```
[2004/04/27 19:25:18, 0, pid=26428] lib/util_sock.c:write_socket_data(407)
write_socket_data: write failure. Error = Connection reset by peer
```

This error occurs in the following cases:

When one CIFS client attempts to access multiple files

When a single CIFS client accesses multiple files, opening or closing each file will require an extended period of time due to virus scans of the files. This forces the subsequent CIFS access requests to have to wait (sometimes, for a long time) causing a session disconnection due to the CIFS client timeout.

When a CIFS client attempts to access a large file

When a CIFS client accesses a large file, opening or closing the file will require more time than usual, due to a virus scan of the file. As a result, the CIFS access request will take long time, and eventually, the CIFS client might time out and the session will be disconnected.

The following describes preventive measures that can be applied if this problem occurs:

When one CIFS client attempts to access multiple files

Change the CIFS client's operation scheme so that, for example, files are accessed sequentially, so as to prevent access to multiple files and to reduce the wait time for virus checking.

When a CIFS client attempts to access a large file

The virus scan for the accessed file will last until it finishes. Therefore, in the event of a previous CIFS access request having timed out, if you try

to access the file again, the virus scan will not be performed again, and you access the file.

When multiple CIFS clients perform CIFS access at the same time, the CIFS clients might experience session disconnections due to timeout, depending on the processing power and the number of the scan servers or the network environment. In this case, you can reduce the wait time caused by virus checking by increasing the number of virus scanning servers.

If you are using Trend Micro ServerProtect and are restricting client host access to nodes by using **Host access restrictions** in the configuration definition for the CIFS service, set access permissions for the host name or network address of the scan server.

The failover functionality might cause virus scans to fail. In this case, the scan will be re-executed when the affected files are accessed again.

If you are using Trend Micro ServerProtect, the number of registered scan servers is included both in the number of logged-in CIFS clients displayed for **Current number of CIFS login clients** on the **CIFS Service Maintenance** page of the **Access Protocol Configuration** dialog box, and in the number of current sessions of MIB information.

Setting home drives

Directories in a CIFS share provided by an HDI system can be used for the home drives of CIFS clients.

Depending on how the settings are configured, a home directory can be automatically created when a home drive is set. If a home directory is not automatically created, you can either create one manually or use a function to automatically create a home directory, which is provided by the HDI system. An example of setting up a home drive is shown below.

To configure settings for individual users by using the Windows Properties window:

The user can specify paths to the home drive (connected drive) and home directory (home folder) from the Properties window. A user set as a CIFS administrator or a user who belongs to a group set as a CIFS administrator must perform this step. Execute this operation in an environment that uses user mapping.

If you use the Windows Properties window to configure the home drive, a home directory will be automatically created.

To configure settings for multiple users at once by using the user registration commands provided by Windows:

When using the command line to register multiple users, you can specify the home drive and home directory paths. If you use the command line to specify the home drive, no home directory will be automatically created.

What is the function for automatically creating a home directory?

On an HDI system, by enabling the function for automatically creating a home directory when creating a shared directory, you can ensure that a home directory is automatically created whenever a CIFS client accesses a CIFS share. The name of the automatically created directory is the user name of the CIFS client in all lowercase letters. Note that the user name of the CIFS client is used as the home directory name if a mixture of uppercase and lowercase letters is allowed in the home directory name on the home-directory-roaming file system. The `cifsoptlist` command can be used to check if a mixture of uppercase and lowercase letters is allowed in the home directory name on the home-directory-roaming file system.

The following figure shows the structure of an automatically created home directory.

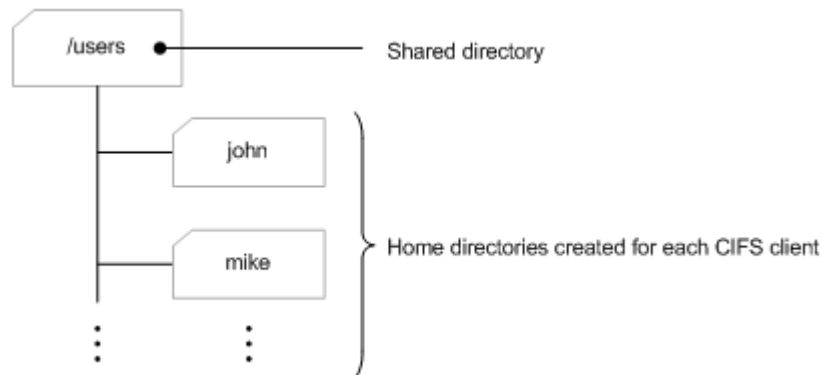


Figure 7-1 Structure of an automatically created home directory

The access permissions for a directory that is automatically created in a file system of the Advanced ACL type depend on the ACL of the parent directory. If there are no ACLs to be inherited from the parent directory, full control access permissions (ACLs are applied to the folder, subfolders, and files) are granted to a CIFS client that uses the home directory. If there are ACLs to be inherited from the parent directory, only those ACLs are applied to automatically generated directories. The access permissions of a CIFS client that uses the home directory are not granted automatically on an individual basis.

The access permissions for a directory that is automatically created in a file system of the Classic ACL type are set as follows:

- For the CIFS client that uses the home directory: `rxwx`
- Group to which the CIFS client using the home directory belongs: `--x`
- All other users: `--x`

Before using the function for automatically creating a home directory

When you are creating a CIFS share, the system administrator must select whether to enable the function for automatically creating a home directory. In making this decision, the following points must be kept in mind:

- A home directory cannot be created when the CIFS client accesses a file system using the guest account (`nobody`).
- If the function for automatically creating a home directory is enabled, a directory will be created even if the CIFS client accessing the file system does not have a home drive set up. The CIFS administrator must delete all unwanted directories.
- If the function for automatically creating a home directory is enabled, a directory will be created even when the file system is accessed by using a computer account. The CIFS administrator must delete all unwanted directories. Alternatively, use an ACL or share-level ACL to restrict access from a computer account so that directories are not created.
- The CIFS client is not notified if an attempt to automatically create a home directory fails. CIFS clients must be instructed to bring up a command prompt and confirm that the home drive was successfully set up.
- If the user name of the CIFS client contains any characters other than the following, the directory must be created manually:
 - Alphanumeric characters
 - Multi-byte characters
 - An exclamation mark (!), a hash mark (#), a dollar sign (\$), a percent sign (%), an ampersand (&), a single quotation mark ('), a left parenthesis ((), a right parenthesis ()), a hyphen (-), a period (.), a caret (^), an underscore (_), a grave accent mark (`), a left curly bracket ({), a right curly bracket (}), a tilde (~), or a space
- If the same user name is used for users in different domains, the automatic creation of directory fails. To prevent the name conflict of the directories, we recommend that CIFS shares be divided by domain and used separately by different domains.
- A user name reserved by the system or any of the following names cannot be used for the user name of a CIFS client. For details about reserved user names, see the *Administrator's Guide*.
 - `.arc`
 - `.backupdates`
 - `.conflict`
 - `.conflict_longpath`
 - `.history`
 - `.lost+found`
 - `.snaps`
 - `.system_gi`
 - `.system_reorganize`
 - `.temp_backupdates`
 - `lost+found`
 - `schedule_syslu_backup.tgz`

- When you use a path name in UNC format (`\\server-name\share-name\...`) to access the home directory, you must not omit the share name.

Using home drives

To start using a home drive, the necessary settings must first be configured in the HDI system, and then the home drive must be registered in the CIFS client environment. Below is a sample procedure for configuring the HDI system, and recommended values for the configuration parameters.

1. Select a CIFS administrator.
In the **CIFS Service Management** page (**Setting Type: Administration**) of the **Access Protocol Configuration** dialog box, as a CIFS administrator, set a user that will use the Windows Properties window to set up the home drives or set a group to which the user belongs.
2. Create and mount a file system.
In the **Create File System** dialog box, create and mount a file system.
3. Create a CIFS share.
Fill in all of the necessary items in the **Add Share** dialog box and create the CIFS share that will serve as the parent directory of the home directory. To prevent improper access to the home directory from other CIFS clients, the following parameter settings are recommended.

Table 7-4 Recommended settings for parameters in the Add Share dialog box

Tab	Item	Recommended setting
Basic	Protocol	Select CIFS (for Windows(R) clients) to specify that the CIFS protocol be used.
	Export point owner	Create directory / change directory owner Select this item to create the shared directory that will become the parent directory of the home directory. Set up an owner user and an owner group as follows for the shared directory to be created: Owner user: <code>root</code> Owner group: <code>root</code>
Access Control	Browsable share	Clear the check box to configure the CIFS client environment to remove the CIFS share name from the list display.
	ACL registered users and groups (When the file system is using the Advanced ACL type.)	Set the ACL for the newly created CIFS share as follows: User or group name Specify <code>Everyone</code> , which is a Windows domain built-in account, for the group. Permissions

Tab	Item	Recommended setting
		Specify access permissions in Full control .
	Access permissions for new directories (When the file system is using the Classic ACL type.)	Set the access permissions for the newly created CIFS share as follows: Owner: RW (allows both read and write) Group: RO (allows read only) Other: RO (allows read only)
Advanced	Enable auto creation of home directory	Select the check box to enable the function that automatically creates a home directory.

Note:

If you use the **Create and Share File System** dialog box to create a CIFS share, in addition to the recommended settings described in the above table, make sure that the **Apply these ACLs to this folder, sub-folders, and files** check box is cleared.

- If necessary, set a default ACL.
To change the access permissions of the home directory, which is automatically created in a file system using the Classic ACL type, use the `dirsetacl` command to set a default ACL.

Notes on using the Windows roaming user profile functionality

This subsection has notes on using the Windows roaming user profile functionality to specify a CIFS share on an HDI system as the destination where the user profiles of CIFS clients are saved.

- Users whose user names include percent signs (%) cannot use the roaming user profile functionality.
- When using the roaming user profile functionality, user profile data is downloaded from a CIFS share of an HDI system when a user logs on to Windows, and then the profile data is applied to the CIFS client. Therefore, if there is a large amount of user profile data, Windows logon processing takes more time. If this is the case, use the folder redirect functionality and configure the settings so that some of the user profile folders (for example, the documents folder, in which a large amount of data might be stored) is redirected to folders under the destination where the user profile is saved.

Files and Folders in a CIFS Share

This chapter contains notes on files and folders that are created in a CIFS shared directory.

- [About file and directory names](#)
- [Owner or group that owns a file or directory](#)
- [Access Control Lists](#)
- [File attributes](#)
- [Timestamps](#)
- [Displaying disk capacity](#)
- [WORM files](#)
- [Access Control by using ABE](#)
- [Restrictions on files and folders on CIFS shares](#)

About file and directory names

This section contains notes on file and directory names.

Supported characters

The HDI system encodes the names of files and directories in UTF-8. The following table describes the maximum length of file and directory names on a CIFS share.

Table 8-1 Maximum length of file and directory names

Item	Maximum length [#]
	Windows/ HDI system
File name	255 characters
Directory name	244 characters
File path name	259 characters
Directory path name	247 characters

#

This is the maximum length of a name or path that can be specified for a file or directory to be accessed from a CIFS client via Explorer. The maximum length might vary depending on the application to be used.

Notes:

- Multi-byte characters are counted as single characters.
- Surrogate pairs are counted as two characters.
- A character to which a code referred to as a variation selector is added is counted as three characters. However, when the character is used in the GUI or in a command of HDI, it is counted as two characters.

When creating or renaming a file or directory, if you specify a name that exceeds the maximum length shown in the above table, an error might occur. When specifying the names of files or directories, specify names that are safely below the maximum length.

Please do not use following characters and path name as the file name or directory name.

- Space at the ending of the file name and/or directory
- Period at the ending of the file name and/or directory (.)
- Double quotation marks ("), asterisk (*), forward slashes (/), colons (:), left angle brackets (<), right angle brackets (>), question marks (?), backslashes (\) and vertical bars (|) in the file name and/or directory

Notes on the maximum lengths of file names and directory names

Accessing files or directories from a CIFS client

Depending on the function used by HDI, a path might be automatically added to the file name or the directory name, causing the file path or the directory path to be longer than the original path. When this happens, if the length of the path exceeds the maximum length, CIFS clients will no longer be able to access the file or directory. Note that the path length of a file or directory contains a host name (or IP address) and a CIFS share name. Read the additional notes below, and make sure that the length of the path of each file or directory does not exceed the maximum. For details about the maximum length of a file or directory path name, see subsection [Supported characters on page 8-2](#).

When using content-sharing functionality :

- If you are using a past version of a directory, the path lengths of files will increase by 25 characters, because the past versions of files migrated to the HCP system are stored in `\.history\YYYY_MM_DD_hhmm\file-path`.
- When the `hcoporphanrestore` command is executed, the files that exist only in the HCP system are restored as follows.

If data is migrated for each file system:

```
/mnt/file-system-name/.lost+found/43-character-HCP-UUID/file-name
```

If data is migrated for each share:

```
/mnt/file-system-name/share-directory-name/.lost+found/43-character-HCP-UUID/file-name
```

- If you use the `arcrestore` command to change the namespace allocation unit from file systems to shares, the data in the file systems that were linked to the HCP system in units of file systems will be migrated to directories in a file system that is set up to link to HCP in units of shares. As a result, path lengths will increase by the number of characters in the name of the migration-destination directory.
- If you use the home-directory roaming functionality, conflict files are stored in `.conflict\YYYY_MM_DD_hhmm_host-name-or-cluster-name@zone-name\file-path`. As a result, path lengths will increase by the number of characters in the string `\.conflict\YYYY_MM_DD_hhmm_host-name-or-cluster-name@zone-name`.
- If you use the read-write-content-sharing functionality and the number of characters in the host name or the share name differs from the corresponding value at some other site that is synchronized with the current site, path lengths might increase.
- If you use the read-write-content-sharing functionality, conflict files are stored in `.conflict\YYYYMMDD\file-path_YYYYMMDDhhmmss`. As a result, path lengths will increase by 34 characters.

When the importing data from another file server :

- When you import data from another file server, the sum of the number of characters in the host name (or IP address) and the CIFS share name at the migration destination might exceed the sum of those values at the migration source. In such cases, the path length will increase by the same number of characters as the difference between the sums.

When using NDMP functionality :

- If you are using the NDMP function and you restore data by specifying a directory whose path name is longer than the original path name, path lengths will increase by the number of characters in the difference between the original path and the specified path.

Accessing files or directories from linked functions via a CIFS client

If the path of a file or directory to be accessed by a CIFS client is too long, a linked function might not work properly. Note that the path length of a file or directory contains a host name (or IP address) and a CIFS share name. Read the additional notes below, and make sure that the length of the path of each file or directory does not exceed the maximum. For details about the maximum length of a file or directory path name, see subsection [Supported characters on page 8-2](#).

- If you are using a Trend Micro scan server, the scan server accesses the files to be scanned as a CIFS client by using the following format:
`\\host-name-or-IP-address\C$\file-system-name\file-path`

MS-DOS file names in 8.3 format

The HDI system generates MS-DOS file names in 8.3 format, which are properly displayed in some applications. However, the naming rules that the HDI system uses differ from the naming rules that Windows uses. In HDI systems, you can make sure that the 8.3-format version of a long file name is valid by executing the following command from the command prompt:

```
dir /x file-name-or-folder-name
```

When multi-byte characters are included in a folder name or file name, the 8.3 file name might be longer than the actual name. For this reason, although the actual folder name or directory name does not reach the maximum path length, the 8.3 file name might. When you use an application program that uses 8.3 file names, be careful that the maximum path length for neither the actual name nor the 8.3 file name is exceeded.

Notes concerning display of a CIFS share name

Some clients display CIFS share names that contain uppercase letters as all lowercase.

For the details of allowed characters for CIFS share name, please refer to the Create and Edit CIFS share items in the *Administrator's Guide*.

Owner or group that owns a file or directory

In an HDI system, Windows built-in users or groups other than `Everyone`, `CREATOR GROUP`, and `CREATOR OWNER` are not recognized.[#] Do not specify such users or groups as the owner or group that owns a file or directory in a CIFS share. In addition, you cannot specify a user whose name begins with an at mark (`@`), or a user who belongs to a domain whose name begins with an at mark.

#

If you used the `cifsoptset` command to enable ACL for the default Windows-domain user groups `Authenticated Users` and `Network`, `Authenticated Users` and `Network` can be recognized. For details on how to enable or disable ACL for the default Windows-domain user groups `Authenticated Users` and `Network`, see the *CLI Administrator's Guide*.

Access Control Lists

The definition of an operation available (or unavailable) to a specific user or group is called an access control entry (ACE). A collection of ACEs is called a discretionary access control list (DACL). The access control list (ACL) supported by NTFS in Windows collectively refers to a DACL, a system access control list (SACL), which is used when a failed or successful resource access attempt is logged, and ACEs related to owners. Note that an ACL sometimes refers to a DACL.

The HDI system supports the following two types of ACL functionality: *Classic ACL* and *Advanced ACL*. Classic ACL can be used to set ACLs that conform to POSIX ACLs. Advanced ACL can be used to set an ACL that conforms to Windows NTFS ACLs. Note, however, that even the Advanced ACL specifications are not completely identical to the Windows NTFS ACL specifications.

This section describes the procedure for specifying ACL settings from a CIFS client, and the ACL and HDI system specifications in Windows. This section also provides notes on using the ACL.

Note that in an HDI system, ACL cannot be set up for the following users or groups:

- A user whose name begins with an at mark (`@`)
- A user who belongs to a domain whose name begins with an at mark (`@`)
- Windows built-in users or groups other than `Everyone`, `CREATOR GROUP`, and `CREATOR OWNER`[#]

#

If you used the `cifsoptset` command to enable ACL for the default Windows-domain user groups `Authenticated Users` and `Network`, `Authenticated Users` and `Network` can be recognized. For details on how to enable or disable ACL for the default Windows-domain user

groups `Authenticated Users` and `Network`, see the *CLI Administrator's Guide*.

Differences between Classic ACLs and Advanced ACLs

The following table shows the applicable range of NTFS ACL items in the HDI system for each type of file system.

Table 8-2 Applicable range of NTFS ACL items in the HDI system

Category	Subcategory#	Classic ACL	Advanced ACL
DACL	Access permission	Three types of permission (<code>rwx</code>) are used.	14 types of permission can be used for fine-grained access control.
	Number of entries that can be set	63 files/126 folders	700 files/700 folders
	Accessing permission	All users	File owner and users that have <code>READ_DAC</code> permission
	Updating permission	Owner and users that have write permission	File owner and users that have <code>WRITE_DAC</code> permission
SACL		Not supported	Not supported
Owner	User	Possible	Possible
	Group	Impossible	Possible
	Changing the owner	Not supported	Possible
	Privileges	Setting ACLs and updating timestamps (POSIX conformance)	Setting and retrieving ACLs, and viewing ownership
	Accessing permission	All users	File owner and users that have <code>READ_DAC</code> permission
	Updating permission	--	Users that have <code>WRITE_OWNER</code> permission
File attributes	Read-only attribute	Settable	Settable
	Archive attribute	Not supported	Settable
	System attribute	Not supported	Settable
	Hidden attribute	Not supported	Settable
	Directory attribute	Settable	Settable

Category	Subcategory#	Classic ACL	Advanced ACL
	Encryption attribute	Not settable	Not settable
	Compression attribute	Not settable	Not settable
	Offline attribute	Settable	Settable
	Normal file attribute	Settable	Settable
	Temporary file attribute	Not settable	Not settable
	Sparse file attribute	Not settable	Not settable
	Reparse point attribute	Not settable	Not settable
	Non-indexed file attribute	Not settable	Not settable
Extended attributes		Not settable	Not settable
File timestamp	Resolution	Second(s)	Second(s)
	Updating permission	Owner and the user having write permission	Users that have WRITE_ATTRIBUTES permission

Legend:

"Possible" indicates that the subcategory can be set. "Impossible" indicates that the subcategory cannot be set. "--" indicates that the subcategory is not applicable.

#

CIFS administrators (root users) registered in File Services Manager are not subject to access control.

Classic ACL type of file system

This subsection provides notes on using the Classic ACL type of file system.

When an ACL of the Classic ACL type is specified by using the `dirsetacl` command, the information displayed for **Name** and **Apply to** for a CIFS client will change depending on the type and setup target of the ACL that has been specified.

The following table displays the relationship between the type and setup target of the ACL that has been specified, and the contents displayed for access control in CIFS clients.

Table 8-3 Relationship between the specified ACL and the contents displayed for access control in CIFS clients

Contents set in the HDI system		Contents displayed in the CIFS client	
ACL type	Setup target	Contents displayed for Name	Contents displayed for Apply to
Access ACL	owner	<i>owner-name</i>	This folder only
	owner group	<i>owner-group-name</i>	
	other	Everyone	
	specific user	<i>user-name or comment-set-during-user-registration</i>	
	specific group	<i>group-name</i>	
	mask	--	--
Default ACL	owner	CREATOR OWNER	Subfolders and files only
	owner group	CREATOR GROUP	
	other	Everyone	
	specific user	<i>user-name or comment-set-during-user-registration</i>	
	specific group	<i>group-name</i>	
	mask	--	--

Legend:

--: Nothing is displayed.

Note:

When the same permissions are set for the access ACL and default ACL for a specific user, specific group, or other, **This folder, subfolders and files** is displayed for **Apply to**.

For the ACL set by the `dirsetacl` command, the information displayed as access permissions on a CIFS client changes depending on the specified permission.

The following table shows the relationship between the specified permissions and the displayed contents of the access permissions for CIFS clients.

Table 8-4 Relationship between the specified permissions and the displayed contents of the access permissions for CIFS clients

Details about access permissions displayed for CIFS clients	Permissions set							
	7 rwx	6 rw-	5 r-x	4 r--	3 -wx	2 -w-	1 --x	0 ---
Traverse Folder/Execute File	A	--	A	--	A	--	A	--
List Folder/Read Data	A	A	A	A	--	--	--	--
Read Attributes	A	A	A	A	--	--	--	--
Read Extended Attributes	A	A	A	A	--	--	--	--
Create Files/Write Data	A	A	--	--	A	A	--	--
Create Folders/Append Data	A	A	--	--	A	A	--	--
Write Attributes	A	A	--	--	A	A	--	--
Write Extended Attributes	A	A	--	--	A	A	--	--
Delete Subfolders and Files	A	--	--	--	--	--	--	--
Delete	A	--	--	--	--	--	--	--
Read Permissions	A	A	A	A	A	A	A	--
Change Permissions	A	--	--	--	--	--	--	--
Take Ownership	A	--	--	--	--	--	--	--

Legend:

A: Allow is set. **--: Allow** is not set.

Procedure for specifying ACL settings from a CIFS client

This subsection describes how to specify ACL settings from a CIFS client.

Specifying ACL settings from a CIFS client

In Windows, the Properties dialog box for a file or folder on an NTFS-formatted disk includes the **Security** item. This item allows you to specify access permissions for each user or group existing in the system or domain. In the HDI system, the ACL settings can be changed only by using this Properties dialog box. The `CACLS` command cannot be used to specify the ACL settings.

Users who can specify ACL settings

In the HDI system, only the file owner and CIFS administrators registered in File Services Manager can set access permissions.

Users who are granted access by configuring the ACL settings

The way the HDI system and Windows handle the ACL access control (under the ACL settings) is different. In Windows, the permission settings made for a group or Everyone affect the owner's permissions. In an HDI system, however, they do not. For example, in Windows, the owner of a file can access that file when permission is set for Everyone. In an HDI system, however, the file owner can access the file only when permission is set for the owner. This is also true for groups.

ACL entry with all access permissions set to blank

In Windows, if all access permissions are set to blank (which means that neither "Allow" nor "Deny" is set) in an ACL entry, the entry itself is deleted. Therefore, if all access permissions are set to blank in the HDI system, the following situations might occur:

- When a file is updated with Microsoft Word, Excel, or PowerPoint, either the entry is deleted or other user's permissions are granted.
- When the access permissions of the owner or owner group are set to blank and an ACL is set in the **Security** tab in the Properties dialog box of the file, the owner is changed to another user.

Accordingly, do not set all access permissions to blank except **Everyone**.

Deny access permission setting

The **Deny** check boxes for access permissions cannot be used for the files and folders in CIFS shares. When you specify ACL settings, use the **Allow** check boxes.

The ACL settings that restrict all types of access

Since ACL settings can only be specified by using **Allow** check boxes, when you want to use an ACL to restrict all types of access for a specific folder or file, set Everyone to None and grant permissions to a specific user or group.

The main ACL operations that delete all access permissions

You cannot delete all the ACL entries or all access permissions that have been set. Attempting to do so has no effect. In this case, set all the access permissions to None. The following shows the main ACL operations that can set all access permissions to None:

- When all access permissions have been set to None for all the files and folders that ACLs are being used for:
Clear all of the **Allow** check boxes for the existing ACLs, and then click **Apply**.
- When all access permissions have been set to None for all the files and folders in a folder that an ACL is being used for:
 - For a folder that is not using the default ACL, select the **Replace all child object permissions with inheritable permissions from this object** check box, and then click **Apply** button.

- If all access permissions are inherited from the parent folder, or if all access permissions for ACL entry that are not inherited from the parent folder are set to None, clear all the check boxes in the **Allow** column of the parent folder's default ACLs, and then click **Apply** button.

How to specify or view the ACL settings for a file

This subsection describes how to specify or view the ACL settings for a file.

Pages for specifying or viewing the ACL settings for a file

The access permission settings for a file can be specified using either the basic settings window or the advanced settings window, as shown in the figure below. To display the basic settings window, open the Properties dialog box for the file, click the **Security** tab, and then click the **Edit** button. To display the advanced settings window, open the Properties dialog box for the file, click the **Security** tab, click the **Advanced** button, click **Change Permissions**, and then double-click the permission entry you want to change.




Figure 8-1 File ACL settings windows (left: basic settings window, right: advanced settings window)

Notes on setting and viewing the ACL for a file

- o Under **Group or user names** in the Properties window, any comments that a system administrator makes about a user, when that user is registered, are displayed.
- o The maximum number of ACLs is 63, and includes the following: an owner (owner), a group (group), everyone (the others), the user (user) registered in the CIFS environment, and a group of virtually mapped Windows groups.
- o The owner of a file cannot be changed.

- The user who created the file, or the group to which that user belongs, cannot be deleted from the ACL.
- The owner's read permission cannot be deleted from the ACL.
- This note applies when you set an ACL for a file from the **Security** tab displayed by choosing **Properties**. In this case, if you attempt to specify an ACL that does not include a combination of a user and a primary group, the ACEs for the owner and owner group of the file are added to the ACL.
- Setting will be ignored even if the execution permission is given to files in CIFS shares.
- The ACL for a file can be set up only by CIFS administrators or the user who created the file in CIFS share.

Notes on file owners

- If a Microsoft Excel, Word, or PowerPoint file matches the following condition, and if the file is updated by using Microsoft Excel, Word, or PowerPoint, then the owner of the file may be changed to another user who is neither the owner of the file before updating nor the user who updates the file.
 - The file before updating has more than one pair of ACL entries, which consists of a user and its primary group.
- If both of the following conditions exist when an ACL is set in the **Security** tab displayed by choosing **Properties** for a file, the file owner changes to a user specified in the ACL:
 - Neither the file owner nor owner group is specified in the ACL.
 - A combination of users who are not the file owner, and their primary group is specified in the ACL.

How to specify or view the ACL settings for a folder

This subsection describes how to specify or view the ACL settings for a folder.

Windows for specifying or viewing the ACL settings for a folder

When you set access permissions for a folder, do not use the basic settings window that is displayed from the **Security** page used to view properties. Instead, use the advanced settings window. To display the advanced settings window, open the Properties dialog box for the folder, click the **Security** tab, click the **Advanced** button, click **Change Permissions**, and then double-click the access permission entry you want to change. The following figure shows the ACL settings window for a folder.

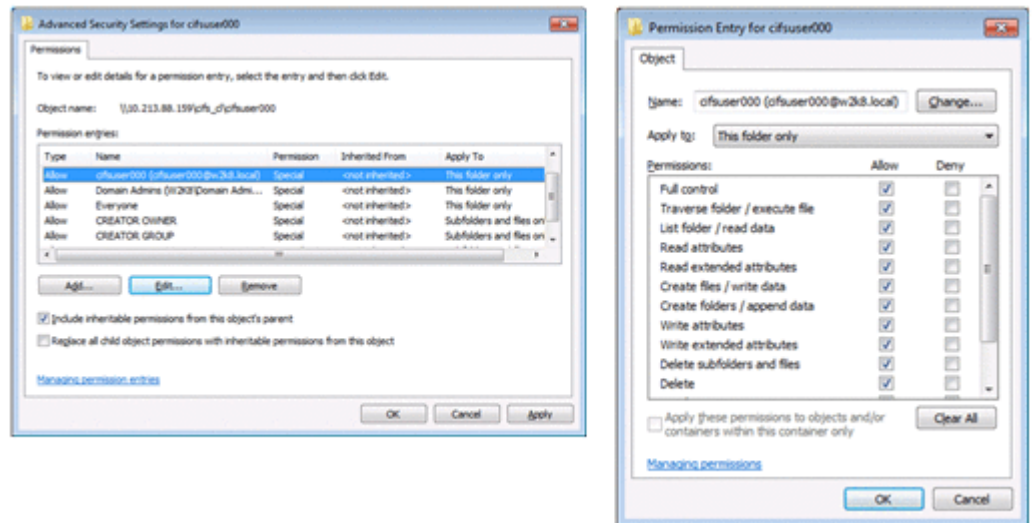


Figure 8-2 Folder ACL settings window

Access ACL and default ACL

A folder has both a default ACL and an access ACL.

The default ACL for a folder controls access to the folder, including the subfolders and files contained in that folder. The access ACL for a folder controls only access to the folder.

In the HDI system, the default ACL for the owner is mapped to Windows `CREATOR OWNER` and the default ACL for the group is mapped to Windows `CREATOR GROUP` when these ACLs are displayed. These ACLs are created when the following operations are performed:

- After a new folder is created in a folder to which the default ACL, user ACLs and group ACLs are added.
- A folder is created in a folder that contains the default ACL.

Procedure for changing ACLs

[Figure 8-3 An example of access permission entries for a folder on page 8-14](#) shows an example of the access permission entries for a folder.

If you want to change only the access ACL, change the ACL for which "This folder only" is displayed as the target in the advanced settings window.

If you want to change the default ACL, the procedure differs depending on whether the target is the owner (or the group to which the owner belongs) of the file, as described below.

- When the target is the owner or the group to which the owner belongs
To set the default ACL for the owner or the group to which the owner belongs, change either `CREATOR OWNER` or `CREATOR GROUP`.
- When the target is neither the owner nor the group to which the owner belongs
The access ACL entries are mapped with names for which **Only this folder** is displayed in the **Apply To** column. The default ACL entries

are mapped with names for which **Subfolders and files only** is displayed in the **Apply To** column. If you change the target to **This folder, subfolders and files**, the same permissions are set for both the default ACL and access ACL.

Changing the access-allowed entry for CREATOR OWNER, CREATOR GROUP, a user other than the owner, or a group other than the group to which the owner belongs when the permission settings are applied to **Subfolders and files only** or **This folder, subfolders, and files** propagates the changes to lower-level folders and files. For details on inheriting access permissions from a parent folder, see [Inheriting access permissions from the parent folder on page 8-17](#).




Figure 8-3 An example of access permission entries for a folder

ACL targets

You can use the advanced ACL view ([Figure 8-2 Folder ACL settings window on page 8-13](#)) to select targets other than those shown above. However, the changes may be applied to either or neither of the access ACL and default ACL, depending on the target. For details on the mappings of the ACL targets with access and default ACLs, see the following table.

Table 8-5 Mapping of ACL targets with access and default ACLs

ACL target that can be selected in Windows	Type of ACL that will be changed		
	When the target is the owner or the owner's group	When the target is CREATOR OWNER or CREATOR GROUP	When the target is a user other than the owner, a group to which the owner does not belong, or Everyone
Only this folder	Access ACL	No changes will be applied	Access ACL

ACL target that can be selected in Windows	Type of ACL that will be changed		
	When the target is the owner or the owner's group	When the target is CREATOR OWNER or CREATOR GROUP	When the target is a user other than the owner, a group to which the owner does not belong, or Everyone
This folder, subfolders and files	--#	No changes will be applied	Access ACL and default ACL
This folder and subfolders	Access ACL	No changes will be applied	Access ACL
This folder and files	Access ACL	No changes will be applied	Access ACL
Subfolders and files only	--#	Default ACL	Default ACL
Subfolders only	No changes will be applied	No changes will be applied	No changes will be applied
Files only	No changes will be applied	No changes will be applied	No changes will be applied

Legend:

--: Not applicable

#

If the target is the owner or the group to which the owner belongs, do not change the **Apply to** value to **This folder, subfolders and files** or **Subfolders and files only**. If you do so, the permissions for the lower-level folders and files might be changed incorrectly. If you want to change the default ACL for the owner (or the group to which the owner belongs), use the ACL for CREATOR OWNER (or CREATOR GROUP).

Note the differences in the way that the contained subfolders and files inherit access permission inheritance depending on the target. For details, see the following table.

Table 8-6 Access permission inheritance for the target, and the contained subfolders and files

ACL target that can be selected in Windows	Type of ACL that inherits access permissions (Type of affected ACL for subfolders and files ^{#1})	
	When the target is the owner or the owner's group	Other cases
Only this folder	No ACL inherits permissions	No ACL inherits permissions

ACL target that can be selected in Windows	Type of ACL that inherits access permissions (Type of affected ACL for subfolders and files ^{#1})	
	When the target is the owner or the owner's group	Other cases
This folder, subfolders and files	No ACL inherits permissions ^{#2}	Access ACL and default ACL for this folder ACL for files
This folder and subfolders	No ACL inherits permissions	Access ACL for this folder
This folder and files	No ACL inherits permissions	ACL for files
Subfolders and files only	No ACL inherits permissions ^{#2}	Access ACL and default ACL for this folder ACL for files
Subfolders only	No ACL inherits permissions	Access ACL for this folder
Files only	No ACL inherits permissions	ACL for files

#1

For the folder and files for which the **Include inheritable permissions from this object's parent** check box is not selected, ACLs are not changed even if this table indicates that they will inherit permissions. For details about this check box, see [Inheriting access permissions from the parent folder on page 8-17](#).

#2

If the target is the owner or the group to which the owner belongs, do not change the "Apply to" value to "This folder, subfolders and files" or "Subfolders and files only." If you do so, the permissions for the lower-level folders and files might be changed incorrectly. If you want to make the permissions for the owner (or the group to which the owner belongs) inheritable, use the ACL for CREATOR OWNER (or CREATOR GROUP).

Notes on specifying and viewing the ACL settings for a folder

- o Information displayed in the **Group or user names** list in the basic ACL view
The **Group or user names** list in the basic ACL view displays the comments that a system administrator has entered during user registration.
- o Maximum number of ACL entries that can be set
Because a folder can contain both the default ACL and access ACL, CREATOR OWNER and CREATOR GROUP are displayed. Therefore, the

maximum number of ACL entries that can be set is 63 for the files times 2 (`CREATOR OWNER` and `CREATOR GROUP`), that is, 126.

- Changing the owner
The owner of a folder cannot be changed.
- The user who created the folder or the group for that user cannot be deleted from the ACLs.
- The owner's access ACL and default ACL (`CREATOR OWNER`) are always set to **Full Control**. These settings cannot be changed. If write permissions are not granted to the owner when a folder is created, **Full Control** is set for the owner's ACLs when the ACL settings are specified for the first time.
- The **Advanced Security Settings** window for a folder includes the **Replace all child object permissions with inheritable permissions from this object** check box ([Figure 8-3 An example of access permission entries for a folder on page 8-14](#)). If you select this check box, the settings for each subfolder and file under the folder are reset to the permissions inherited from the parent directory (default ACL settings for the parent directory). The inheritance of access permissions is also enabled. If, however, the default ACL has not been set for the parent directory, there is no inheritable ACL and, therefore, the ACLs for the subfolders and files under the folder will not be changed.
- If ACL is set in a property window of a folder, full control (`rwx`) will be set on its mask by whatever its access permission.
- The ACL for a file can be set up only by CIFS administrators or the user who created the folder in CIFS share.

Inheriting access permissions from the parent folder

This subsection describes how to have access permissions inherited from the parent folder.

Access permission inheritance for a new file or folder

See the advanced settings window ([Figure 8-2 Folder ACL settings window on page 8-13](#)). In this window, check whether an ACL to be applied to **Subfolders and files only** or **This folder, subfolders and files** exists. When such an ACL exists, if you create a new file or subfolder in a folder, the file or subfolder will inherit the default ACL from the parent folder. The default ACL is applied to all the access ACLs other than the access ACL for the owner, the access ACL for the group to which the owner belongs, and the access ACL for Everyone.

The following table lists the access ACL values that are set when creating a new file or folder in an HDI CIFS share.

Table 8-7 Access ACL values that are set when creating a new file or folder in an HDI CIFS share

Default ACL	Target entry	Access ACL values set for a new file or folder
Not set	Owner	Value of Access permissions for new files or Access permissions for new directories
	Owner group	Value of Access permissions for new files or Access permissions for new directories
	Everyone	Value of Access permissions for new files or Access permissions for new directories
	Added ACEs	Non-existent
Owner Owner group Everyone	Owner	Logical AND of the default ACL and the value of Access permissions for new files or Access permissions for new directories
	Owner group	Logical AND of the default ACL and the value of Access permissions for new files or Access permissions for new directories
	Everyone	Logical AND of the default ACL and the value of Access permissions for new files or Access permissions for new directories
	Added ACEs	Non-existent
Owner Owner group Everyone Added ACEs	Owner	Value of Access permissions for new files or Access permissions for new directories
	Owner group	Value of Access permissions for new files or Access permissions for new directories
	Everyone	Value of Access permissions for new files or Access permissions for new directories
	Added ACEs	Default ACL

Access permission inheritance for an existing file or folder

The **Advanced Security Settings** window for a file or folder includes the **Include inheritable permissions from this object's parent** check box (the following figure). When you select this check box and subsequently change the access permission settings for the parent directory, the changes are automatically applied to the subfolders and files contained in the folder. If you want to set the access permissions on a per-file or per-folder basis, do not select the check box.




Figure 8-4 Inherit access permissions check box

Whether an existing file or folder can inherit access permissions

The **Include inheritable permissions from this object's parent** check box that appears as a file or folder property is selected in the following cases:

- When the default ACL is not set for the parent folder
- When the default ACL is set for the parent folder, and the subfolders and files under the folder have the same default ACL settings as the parent's default ACL settings

To explicitly disable this setting, clear the check box and click **Apply**. Note that only CIFS administrators registered in File Services Manager and the owner of a file or folder can clear the check box.

If the `XCOPY` command or a backup utility is used to migrate resources from a Windows domain environment, the **Include inheritable permissions from this object's parent** setting is also migrated together with the ACL information.

The following table describes the differences in the user's ACL inheritance operation between Windows and the HDI system when the ACL settings for a folder are changed.

Table 8-8 Differences in user operations for inheriting ACL settings

Server	Client
	Windows
Windows	For the subfolders and files, select the Include inheritable permissions from this object's parent check box.
File Services Manager	Make sure that the ACL settings are the same as the default ACL settings of the parent folder. [#]

#

In the HDI system as well, if the default ACL of the parent folder has more permissions than the target access ACL, the target access ACL can inherit the ACL from the parent folder by selecting the following check box: **Include inheritable permissions from this object's parent**. If you do not select the check box, you must manually set the access ACL to have the same settings as the default ACL of the parent folder.

Notes on access permission inheritance for an existing file or folder

- o Even when the **Include inheritable permissions from this object's parent** check box is selected, access permissions are not changed for a folder or file having a different owner from that which changed the access permissions for the parent folder. To change the access permissions for such a folder or file, a CIFS administrator registered in File Services Manager must specify the ACL settings or the owner of the folder or file must directly specify the ACL settings.
- o This note applies when inheritance from the parent is enabled, the default ACL is set for a folder, and the owner and updater of a file belong to different primary groups. Under these conditions, if the file is updated with Microsoft Excel, Word, or PowerPoint, the ACL settings for the pre-update owner might change to the ACL settings for the post-update owner. Similarly, the ACL settings for the pre-update group to which the owner belongs might change to the ACL settings for the post-update group to which the owner belongs. These changes might cause the former owner and the users who belong to the primary group for the former owner cannot access the file.

Adding user ACLs or group ACLs

This subsection describes how to add user or group ACLs.

User or group ACLs can be added by choosing **Add** in the access permission window for a file or folder. The following figure shows the **Select User or Group** window that appears when the **Add** button is clicked.

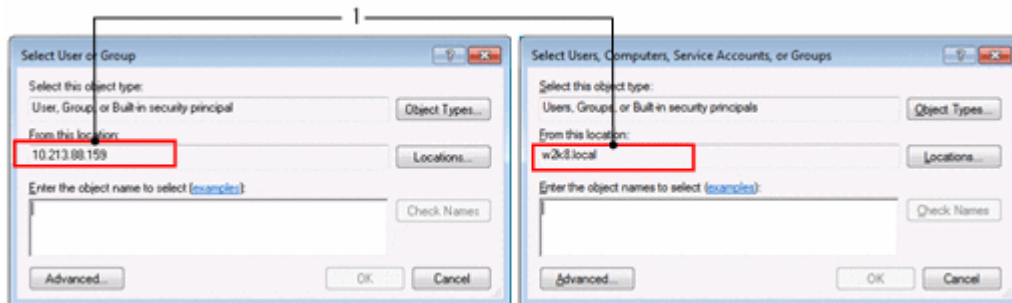


Figure 8-5 Select User or Group window (Left: When user mapping is not being used; Right: When user mapping is being used)

If you specify ACL settings for a folder created in a CIFS file share by using the file or folder properties window, the location of a user or group you select

differs depending on the CIFS service authentication method. The location mentioned above corresponds to the number 1 in the above figure.

When user mapping is not used for local authentication, Active Directory authentication, or NT domain authentication

You must select a user or group for which the host name of an HDI node is displayed in the **From this location** box of the **Select User or Group** window.

Notes:

- When Active Directory authentication or NT domain authentication is selected, a domain controller might be displayed in the **From this location** box. Do not specify it. If it is specified, the ACL setting will not be enabled.
- If you want to display a user (the host name of an HDI node) in the **From this location** box on the left side of [Figure 8-5 Select User or Group window \(Left: When user mapping is not being used; Right: When user mapping is being used\)](#) on page 8-20, register a CIFS user on the File Services Manager **Add User** page of the **Local Users** dialog box.
- To display a group in that box, select the **Apply to CIFS ACL environment** check box when adding a group on the File Services Manager **Add Group** page of the **Local Users** dialog box.
- Depending on the environment, user or group ACLs might not be able to be added in the access permission window for a file or folder. In such cases, use the `dirsetacl` command to specify ACL settings. If you want to add an ACL to a file, use the `dirsetacl` command to specify the ACE inheritance range for the relevant directory so that the required ACEs are inherited by the target file.

When user mapping is used for Active Directory authentication or NT domain authentication

You must select a user or group for which the domain controller is displayed in the **From this location** box in the **Select User or Group** window.

If the domain controller is not displayed for the **From this location** box in the **Select User or Group** window, communication with the domain controller might fail. Possible causes are as follows:

- o An incorrect value is specified for **Domain name (NetBIOS)** in the **Active Directory Authentication** page of the **Access Protocol Configuration** dialog box.
- o On the CIFS client, the IP address of the domain controller cannot be resolved by DNS.
- o The user who is performing CIFS is not a domain user.

The ACL of a local group or local user can be configured by selecting a user or group for which the host name of an HDI node is displayed in the **From this location** box from the **Select User or Group** window.

Notes:

- When you configure the ACL of an HDI local group or local user, use the **Select User or Group** window on the left side of [Figure 8-5 Select User or Group window \(Left: When user mapping is not being used; Right: When user mapping is being used\) on page 8-20](#). A CIFS user must be registered via the **Add User** page of the File Services Manager **Local Users** dialog box to display a user for which the host name of an HDI node is displayed in the **From this location** box.
- To display a group in that box, select the **Apply to CIFS ACL environment** check box when adding a group on the File Services Manager **Add Group** page of the **Local Users** dialog box.
- A CIFS client must be registered in a domain, or no HDI local groups or local users are displayed, even if the host name of an HDI node is displayed in the **From this location** box in the **Select User or Group** window.
- In some environments, you might not be able to add a user or group ACL from the access permission window for a file or folder. In such cases, use the `dirsetacl` command to specify the necessary settings. To add an ACL to a file, use the `dirsetacl` command to set the inheritance range of the access control entries (ACEs) for the relevant directory so that the file inherits the ACEs.

Note that, regardless of the CIFS service authentication format, if you cannot add a user ACL or group ACL from the file or folder access permissions window, you might be able to add the ACL by using one of the following methods:

- Log on to the Windows client as an administrator (a user belonging to the Administrators group).
 - Log on to the Windows client as user who is not an administrator (as a user who was individually added to the Windows client).
 - When the authentication dialog box for adding a user ACL or a group ACL appears, enter your user name and password for connecting to HDI system.
 - Make sure the user name and password you use to connect to HDI system and those you use for logging on to the Windows client are the same.
 - Connect to HDI system by specifying the host name from the Windows client.
 - Make sure the Windows client is a participant in the Active Directory domain.
- If you do this, no problem will occur even if the Windows client has no relationship to the HDI system authentication method or to the domain in which HDI system participates.

ACL set for a newly created file

Because of conformance with POSIX in the HDI system, the owner and group are always added to the ACL at the time of file creation. For details about the

ACL of a new file, see [Table 8-7 Access ACL values that are set when creating a new file or folder in an HDI CIFS share on page 8-18](#).

ACL set for a newly created folder

Same as a file, the owner and group are always added to the ACL for a folder at the time of creation.

For details on the ACL of a new folder, see [Table 8-7 Access ACL values that are set when creating a new file or folder in an HDI CIFS share on page 8-18](#), which also applies to creating new files. For the ACL for the owner, Full Control is always set as an ACL property.

SACL

The request of setting SACL from CIFS clients will be ignored. If setting is requested, the operation will be accepted but actual setting will not be done.

Invalid ACE

The ACEs of BUILTIN/well-known SID accounts and ACEs for which UIDs or GIDs cannot be resolved are ignored, and only other ACEs are set. This is also the case when UIDs and GIDs registered in the Active Directory or the LDAP server have not been mapped.

Mapping ACL specifications in Windows to file permissions in the HDI system

In the HDI system, to provide a POSIX-compliant ACL, file permissions (rwx) in Linux are mapped with the items shown in the basic and advanced ACL views. The following table describes the relationship between the Windows access permissions items displayed on a CIFS client and the file permissions in the HDI system.

Table 8-9 Relationship between the Windows access permissions items and the file permissions in the HDI system

#	Windows access permissions items		File permissions in the HDI system
	Basic	Advanced	
1	Read	Folder List / Read Data	r - -
2		Read Attributes	
3		Read Extended Attributes	
4	Read and Execute	Items 1-3 and 11	r - x
5	Write	Create Files / Write Data	- w -
6		Create Folders / Append Data	
7		Write Attributes	
8		Write Extended Attributes	
9	Change	All Allow check boxes checked	r w x

#	Windows access permissions items		File permissions in the HDI system
	Basic	Advanced	
10	Full Control	All Allow check boxes checked	r w x
11	--	Traverse Folders / Execute File	- - x#1
12		Delete Subfolders and Files#2	- - -
13		Delete	
14		Read Permissions	
15		Change Permissions	

Legend:

--: No setting exists in the basic ACL views.

#1

If an executable file is stored in a CIFS share of the HDI system, even though you do not have Execute File permission, you can execute the file with only Read permission for the file.

#2

In the HDI system, the Delete Subfolders and Files permission is included in the Write permission. That is, if you have Write permission on a folder, you can delete the files and subfolders in the folder.

Advanced ACL type of file system

This subsection provides notes on using the Advanced ACL type of file system.

Setting and displaying an ACL from a CIFS client

The following describes ACL settings and how to display an ACL from a CIFS client.

In a file system of the Advanced ACL type, only the account or CIFS administrator who are allowed to read access permissions and change ownership in the File or Folders Properties window for the file or folder created in a CIFS share, or who has been registered in the CIFS service, can set up an ACL.

Access permissions that can be specified for a file or directory

The following table shows the access permissions that can be set for the Advanced ACL type of file system from a CIFS client and the applicable NTFS ACE masks. For each item in the table, you can choose to allow or deny the given operation. If both allowing and denying an operation are specified, the deny specification takes precedence.

Table 8-10 Access permissions that can be specified in an ACL and the applicable NTFS ACE masks

#	Access permission	Operation permitted or prohibited	NTFS ACE masks
1	Traverse Folder# ¹	Allows or denies moving through folders to reach other files or folders even if the user has no permissions for the traversed folders.	FILE_TRAVERSE
	Execute File# ²	Allows or denies running program files.	FILE_EXECUTE
2	List Folder# ¹	Allows or denies viewing file names and subfolder names within the folder.	FILE_LIST_DIRECTORY
	Read Data# ²	Allows or denies viewing data in files.	FILE_READ_DATA
3	Read Attributes	Allows or denies viewing attributes of a file or folder, such as read-only and hidden.	FILE_READ_ATTRIBUTES
4	Read Extended Attributes	Allows or denies viewing the extended attributes of a file or folder.	FILE_READ_EA
5	Create Files# ¹	Allows or denies creating files within the folder	FILE_ADD_FILE
	Write Data# ²	Allows or denies making changes to the file or overwriting existing content.	FILE_WRITE_DATA
6	Create Folders# ¹	Allows or denies creating folders within the folder.	FILE_ADD_SUBDIRECTORY
	Append Data# ²	Allows or denies making changes to the end of the file but not changing, deleting, or overwriting existing data.	FILE_APPEND_DATA
7	Write Attributes	Allows or denies changing the attributes of a file or folder such as read-only and hidden.	FILE_WRITE_ATTRIBUTES
8	Write Extended Attributes	Allows or denies changing the extended attributes of a file or folder.	FILE_WRITE_EA
9	Delete Subfolders and Files# ¹	Allows or denies deleting subfolders and files even if the Deleting permission has not been granted on the subfolder or file.	FILE_DELETE_CHILD
10	Delete	Allows or denies deleting the file or folder (Note that if you do not have Delete permission on a file or folder, you can still delete it if you have been granted Delete Subfolders and Files on the parent folder.)	DELETE
11	Read Permissions	Allows or denies reading permissions of the file or folder.	READ_CONTROL

#	Access permission	Operation permitted or prohibited	NTFS ACE masks
12	Change Permissions	Allows or denies changing permissions of the file or folder.	WRITE_DAC
13	Take Ownership	Allows or denies taking ownership of the file or folder.	WRITE_OWNER

#1

Applies to folders only.

#2

Applies to files only.

Setting access permissions for a file or directory

For Advanced ACL, the access permission information sent from a CIFS client is converted to an HDI-specific format and then set to the file system. ACEs are saved in the order in which they are sent from the CIFS client.

If a built-in or well-known SID account is encountered, or if an ACE for which the UID or GID cannot be resolved is encountered, the entry is skipped.

The following figure shows an overview of setting up an ACL:

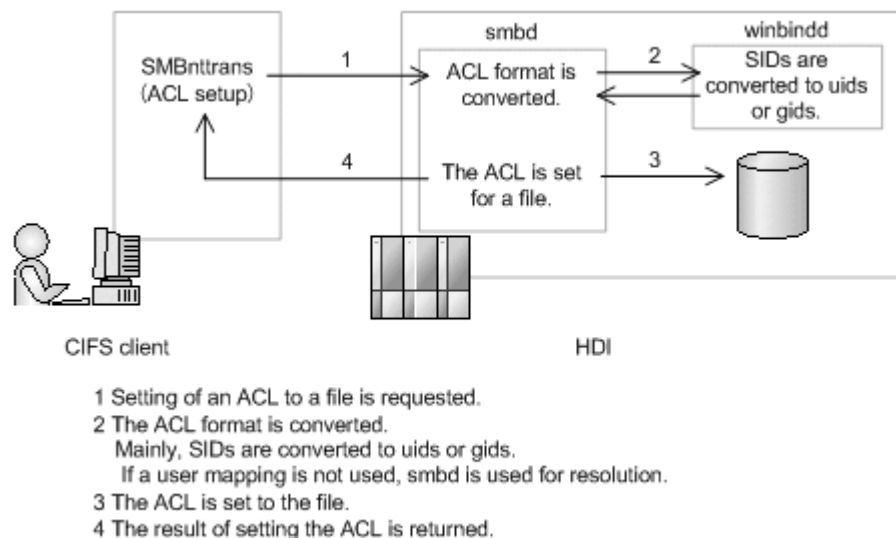



Figure 8-6 Overview of setting up an ACL

Acquiring the access permission information for a file or directory

The access permission information acquired from a file or directory is converted to a Windows-specific format, and then returned to the CIFS client.

The following figure shows an overview of obtaining an ACL:



- 1 Obtaining an ACL is requested.
- 2 Obtaining the ACL from a file.
- 3 The ACL format is converted to a Windows specific format.
Mainly, uids or gids are converted to SIDs.
If user mapping is not used or user mapping cannot be used to resolve SIDs, smbd is used for resolution.
- 4 The ACL is returned.

Figure 8-7 Overview of obtaining an ACL

File system root ACL

In the initial status after a new Advanced ACL type of file system is created, the file system root ACL is set as shown in the following table.

Table 8-11 Default values for the file system root ACL

Name	Access permission	Target	Windows Server 2008, 2012, 2016	Advanced ACL
Administrators	Full Control	This folder, subfolders and files	Y	--
SYSTEM	Full Control	This folder, subfolders and files	Y	--
CREATOR_OWNER	Full Control	Subfolders and files only	Y	--
Users	Read & Execute	This folder, subfolders and files	Y	--
Users	Create Folders / Append Data	This folder and subfolders	Y	--
Users	Create Files / Write data	Subfolders only	Y	--
Everyone	Full Control	This folder, subfolders and files	Y#	Y

Legend:

Y: The file system root ACL exists.

--: The file system root ACL does not exist.

#

The access permission is "Read & Execute". The target is "Only this folder".

ACL-related values

An ACL consists of ACEs that define access permissions for each user or group. Each ACE consists of the following four elements:

- User name or group name (or equivalent ID)
- ACE type, which defines whether the ACE entry allows or denies the operation
- ACE mask, which defines the operation to be permitted or prohibited
- ACE flag, which defines ACL inheritance and other items

The following tables show the ACE type, ACE mask, and ACE flag values used in an NTFS ACL and the support status in Advanced ACL of the HDI system.

Table 8-12 List of ACE types of Advanced ACL

#	ACE type	Explanation	Supported
1	ACCESS_ALLOWED_ACE_TYPE	This ACE is ALLOW entry.	Y
2	ACCESS_ALLOWED_CALLBACK_ACE_TYPE	Applications start callback function when allowed access.	N
3	ACCESS_ALLOWED_CALLBACK_OBJECT_ACE_TYPE	Specialized ACE type for #2 object	N
4	ACCESS_ALLOWED_COMPOUND_ACE_TYPE	(reserved)	N
5	ACCESS_ALLOWED_OBJECT_ACE_TYPE	Allow entry for Active Directory objects	N
6	ACCESS_DENIED_ACE_TYPE	This ACE is DENY entry.	Y
7	ACCESS_DENIED_CALLBACK_ACE_TYPE	Applications start callback function when denied access	N
8	ACCESS_DENIED_CALLBACK_OBJECT_ACE_TYPE	Specialized ACE type for #7 object.	N
9	ACCESS_DENIED_OBJECT_ACE_TYPE	Deny entry for Active Directory objects	N
10	ACCESS_MAX_MS_ACE_TYPE	(reserved)	N
11	ACCESS_MAX_MS_V2_ACE_TYPE	(reserved)	N
12	ACCESS_MAX_MS_V3_ACE_TYPE	(reserved)	N
13	ACCESS_MAX_MS_V4_ACE_TYPE	(reserved)	N

#	ACE type	Explanation	Supported
14	ACCESS_MAX_MS_OBJECT_ACE_TYPE	(reserved)	N
15	ACCESS_MIN_MS_ACE_TYPE	Same with ACCESS_ALLOWED_ACE_TYPE	N
16	ACCESS_MIN_MS_OBJECT_ACE_TYPE	Same with ACCESS_ALLOWED_OBJECT_ACE_TYPE	N
17	SYSTEM_AUDIT_ACE_TYPE	Related to audit	N
18	SYSTEM_ALARM_ACE_TYPE	(reserved)	N
19	SYSTEM_ALARM_CALLBACK_ACE_TYPE	(reserved)	N
20	SYSTEM_ALARM_CALLBACK_OBJECT_ACE_TYPE	(reserved)	N
21	SYSTEM_ALARM_OBJECT_ACE_TYPE	(reserved)	N
22	SYSTEM_AUDIT_CALLBACK_ACE_TYPE	Related to audit	N
23	SYSTEM_AUDIT_CALLBACK_OBJECT_ACE_TYPE	Related to audit	N
24	SYSTEM_AUDIT_OBJECT_ACE_TYPE	Related to audit	N

Legend:

Y: Supported, N: Not supported

Table 8-13 List of NTFS ACE masks and the support status

Bit	Access permission	Name in the Windows GUI	Explanation	Supported
31	GENERIC_READ	--	Appends all of the following flags: FILE_READ_ATTRIBUTES FILE_READ_DATA FILE_READ_EA READ_CONTROL SYNCHRONIZE	Y#1
30	GENERIC_WRITE	--	Appends all of the following flags: FILE_APPEND_DATA FILE_WRITE_ATTRIBUTES FILE_WRITE_DATA FILE_WRITE_EA READ_CONTROL SYNCHRONIZE	Y#1

Bit	Access permission	Name in the Windows GUI	Explanation	Supported
29	GENERIC_EXECUTE	--	Appends all of the following flags: FILE_READ_ATTRIBUTES READ_CONTROL SYNCHRONIZE FILE_EXECUTE	Y#1
28	GENERIC_ALL	--	Appends all of the following flags: DELETE_ACCESS READ_CONTROL_ACCESS WRITE_DAC_ACCESS WRITE_OWNER_ACCESS SYNCHRONIZE_ACCESS FILE_ALL_ACCESS (all bits, from 0 to 8)	Y#1
27	(reserved)	--	--	N#2
26	(reserved)	--	--	N#2
25	(reserved)	--	--	N#2
24	RIGHT_TO_ACCESS_SACL	--	--	N#3
23	(unallocated)	--	--	N#2
22	(unallocated)	--	--	N#2
21	(unallocated)	--	--	N#2
20	SYNCHRONIZE	Synchronize	The right to specify a file handle in one of the wait functions. However, for asynchronous file I/O operations, you should wait on the event handle in an overlapped structure rather than using the file handle for synchronization.	Y
19	WRITE_OWNER	Take Ownership	Takes ownership.	Y
18	WRITE_DAC	Change Permissions	Enables changing of DACL.	Y
17	READ_CONTROL	Read Permissions	Enables reading of DACL.	Y
16	DELETE	Delete	Enables to delete files or directories. Available to do even when FILE_DELETE_CHILD	Y

Bit	Access permission	Name in the Windows GUI	Explanation	Supported
			deny entry is set to parent directories.	
15	(unallocated)	--	--	N#2
14	(unallocated)	--	--	N#2
13	(unallocated)	--	--	N#2
12	(unallocated)	--	--	N#2
11	(unallocated)	--	--	N#2
10	(unallocated)	--	--	N#2
9	(unallocated)	--	--	N#2
8	FILE_WRITE_ATTRIBUTES	Write Attributes	Enables to write NTFS attributes.	Y
7	FILE_READ_ATTRIBUTES	Read Attributes	Enables to read NTFS attributes.	Y
6	FILE_DELETE_CHILD	Delete Subfolders and Files	Enables to delete files and subdirectories in directories. When "read-only attribute" is set on files or subdirectories, an attempt to delete fails.	Y
5	FILE_EXECUTE	Execute File	Enables to execute files.	Y
	FILE_TRAVERSE	Traverse Folder	Enables to traverse directories (for compatibility with UNIX.)	Y
4	FILE_WRITE_EA	Write Extended Attributes	Enables to write extended attributes.	Y#4
3	FILE_READ_EA	Read Extended Attributes	Enables to read extended attributes.	Y#4
2	FILE_APPEND_DATA	Append Data	Enables to append data to files.	Y
	FILE_ADD_SUBDIRECTORY	Create Folders	Enables to create subdirectories in directories.	Y
1	FILE_WRITE_DATA	Write Data	Enables to write data to files.	Y
	FILE_ADD_FILE	Create Files	Enables create files in directories.	Y
0	FILE_READ_DATA	Read Data	Enables to read data in files.	Y
	FILE_LIST_DIRECTORY	List Folder	Enables to list contents in directories.	Y

Legend:

Y: Supported

N: Not supported

--: Not applicable

#1

The `GENERIC_READ`, `GENERIC_WRITE`, `GENERIC_EXECUTE`, and `GENERIC_ALL` do not have individual access rights. They are flags for setting access permissions that are independent from ACE setting objects: files, directories, Active Directory, and so on. Setting these flags creates new multiple access rights.

#2

When it receives currently unallocated bits, they are turned off and the processing continues.

#3

This bit (mask) should not be turned on in an ACE for a file or directory. Also, because the HDI system does not support SACL, if this bit is set in an ACE that is received, an error is returned to the client.

#4

Extended attributes are specific to OS/2 HPFS. Also, XFS does not support extended attributes. As such, it is not necessary for the HDI system to do anything to these bits (masks) either. However, these bits are saved and ensured in the file systems to avoid losing information when copied client files in HDI.

Table 8-14 NTFS ACE flags and supports

bit	ACE flag	Explanation	Support status
7	<code>FAILED_ACCESS_ACE_FLAG</code>	Stores audit message for failed access attempts.	N
6	<code>SUCCESSFUL_ACCESS_ACE_FLAG</code>	Stores audit message for successful access attempts	N
5	--	Unused (use in future as group bits)	--
4	<code>INHERITED_ACE</code>	Indicates that the ACE has been inherited for generations.	Y
3	<code>INHERIT_ONLY_ACE</code>	Indicates an inherit-only ACE, which does not control access to the object to which it is attached. Only subsequent generations can inherit it.	Y
2	<code>NO_PROPAGATE_INHERIT_ACE</code>	Neither of the <code>OBJECT_INHERIT_ACE</code> flag nor the <code>CONTAINER_INHERIT_ACE</code> flag are inherited. Corresponded to "Apply these permissions to objects and	Y

bit	ACE flag	Explanation	Support status
		containers within this container only."	
1	CONTAINER_INHERIT_ACE	Directories inherit the ACE as an effective ACE.	Y
0	OBJECT_INHERIT_ACE	Files inherit the ACE as an effective ACE.	Y

Legend:

Y: Supported

N: Not supported

--: Not applicable

Table 8-15 The range where the permission is applied and states of ACE flag

Where the permission is applied ("Apply To" column in the "Permission entries" table in the "Advanced Security Settings" dialog box)	States of ACE flags [#]		
	bit 3	bit 1	bit 0
This folder only	0	0	0
This folder, subfolders, and files	0	1	1
This folder and subfolders	0	1	0
This folder and files	0	0	1
Subfolders and files only	1	1	1
Subfolders only	1	1	0
Files only	1	0	1

#:

The meaning of each bit of ACE flags is described in [Table 8-14 NTFS ACE flags and supports on page 8-32](#)

ACL evaluation

Whether a file or directory access request is allowed or denied is determined by the following rules:

- If no ACL is set (NULL ACL), all types of access requests are permitted.
- Any access is denied when there is no ACE (Empty ACL). Only file owner can change access rights considered as having `READ_CONTROL` and `WRITE_DAC` flags ON.
- The file owner ACEs, file owner group ACEs, and Everyone ACEs are evaluated.

The file owner ACEs and file owner group ACEs are evaluated at the location where denial to the file or directory occurs. The Everyone ACEs are evaluated after other ACEs, in the order specified in the ACL list.

- ACEs are evaluated in the listed order.
- ACEs are not evaluated after the evaluation result has been settled.
- Deny is settled as ACE's evaluation when the DENY entry is found.
- ALLOW is settled as ACE's evaluation when the ALLOW entry is found.
- Evaluation results in DENY when the whole entries cannot be settled.

Each CIFS client in Windows takes responsibility of ordering ACEs. The CIFS client sorts ACEs in the following order: "deny" ACEs of itself, "allow" ACEs of itself, "deny" ACEs inherited from the parent, "allow" ACEs inherited from the parent, "deny" ACEs inherited from the grandparent, and "allow" ACEs inherited from the grandparent. The CIFS client then requests the CIFS server to store the ACL. The HDI system also expects the client to store the ACEs in the above order, and will not check whether the order of ACEs is correct.

The order of the ACEs under the **Advanced** window (accessed by clicking the **Advanced** button in the **Security** tab of the Properties window) is the same as the ACE's evaluation (the order of the ACL list).

Since ACEs are evaluated as described above, so that great care is required when Everyone has Deny ACE and a user has ALLOW ACE: otherwise, a user cannot access any object.

ACL initial values, inheritance, and propagation

The ACL initial values for newly created files or directories are inherited from their parents.

This ACL inheritance chain can be broken. When you break the ACL inheritance chain of files or directories, you can select whether the ACL settings that have been inherited are discarded or whether the ACL settings that were inherited before the chain was broken are incorporated into the file or directory as its own ACL settings.

After you break the ACL inheritance chain, you can restore it. Note that if you restore the chain for a file or directory in which ACL settings have been incorporated as its own ACL settings, the same ACL settings are set twice.

When an ACL that has the inheritance attribute is changed, the change is propagated to the child and later offspring on the CIFS client side. The change cannot be propagated to a file or directory that is accessed via NFS or another protocol.

Accordingly, the responsibility for propagating changes to ACL settings of a file or directory accessed by clients other than CIFS clients belongs to the application.

ACE duplication check

The HDI system does not check for duplicated ACE entries registered for the same user or group from a CIFS client.

SACL

The request of setting SACL from CIFS clients will be ignored. If setting is requested, the operation will be accepted but actual setting will not be done.

Invalid ACE

The ACEs of BUILTIN/well-known SID accounts and ACEs for which UIDs or GIDs cannot be resolved are ignored, and only other ACEs are set. This is also the case when UIDs and GIDs registered in the Active Directory or the LDAP server have not been mapped.

File owners and UNIX permissions

For a file owner of an Advanced ACL type of file system, either a user or group can be registered. As shown in the following table, the HDI system internally performs mapping among file owners, file owner users (UNIX permissions), and file owner groups (UNIX permissions). Therefore, be careful when you view information or change permission settings for the relevant file from NFS.

Table 8-16 Handling of file owners for UNIX permissions

File owner	Handling for UNIX permissions	
	File owner user	File owner group
User	UID of the file owner	GID of the primary group to which the file owner belongs
Group	"groupowner" (<i>groupowner</i> is the name of the user who is assigned to the UID internally used for the system.)	GID of the file owner

Accounts for setting as owners

In a file system of the Advanced ACL type, only the account who is allowed to read access permissions and take ownership, or CIFS administrator who has been registered in the CIFS service, can change the owner of a file or folder created in a CIFS share.

The following table lists the types of account and indicates whether each account can be set as an owner.

Table 8-17 Whether an account can be set as an owner

Account type	Settable as owner	Owner privilege
User	Yes	The user
Group	Yes	All users belonging to the group
BUILTIN/well-known SID account	No [#]	--

Account type	Settable as owner	Owner privilege
Account whose SID cannot be resolved (such as a user outside the domain or a deleted account)	No#	--

Legend:

--: Not applicable

#

The processing does not return an error but is just skipped. This is a workaround to avoid errors occurring (resulting in the processing being aborted) for data migration scripts using the `XCOPY` command.

Setting the owner for a file or directory

The owner information (SID) sent from a CIFS client is converted to a UID or GID in the HDI system and set for the file system.

If the requested account is a BUILTIN/well-known SID account or an account whose UID or GID cannot be resolved, processing normally terminates with nothing performed.

Acquiring the owner of a file or directory

The owner information (UID or GID) acquired from a file or directory is converted to SID in the HDI system and is then returned to the CIFS client.

If there are entries that cannot be converted into an SID (for example, accounts that are not subject to CIFS management, such as users accessing data via NFS), the HDI system generates an HDI-system-specific SID (this is also true for the Classic ACL type). In this case, the SID instead of the user name is displayed in the CIFS client.

Whether the owner group can be set

In the same way as the Classic ACL type of file system, the Advanced ACL type of file system can store POSIX-compatible owner group information of files/directories.

Unlike the owner, an owner group is not set when a file is created. Also, the owner group cannot be manipulated by the use of ordinary operations (for example, from the Properties window), and can only be manipulated by using Windows commands.

In the Advanced ACL type of file system, owner groups are not used for checking access permission, and are used only for quota management. Also, in the Advanced ACL type of file system, the owner group cannot be set if a group is set as the owner.

The following table shows whether the owner group can be set depending on the owner.

Table 8-18 Whether the owner group can be set

Owner	Account	Owner group settable
User	Group	Yes
	BUILTIN/Well-known SID account	No
	Account whose SID cannot be resolved (such as an extra-domain user or deleted account) [#]	No [#]
Group	--	No [#]

Legend:

--: Not applicable

#

The processing does not return an error but is just skipped. This is a workaround to avoid errors occurring (resulting in the processing being aborted) for data migration scripts using the `XCOPY` command.

Setting the owner group for a file or directory

If a built-in or well-known SID account or an account where the GID cannot be resolved from the SID (such as an extra-domain group or user), processing terminates normally but the owner group is not changed.

When an attempt is made to change the primary group for a file or directory for which a group is set as the owner, processing terminates normally but the owner group is not changed.

Maximum number of ACL entries that can be set

In the Advanced ACL type of file system, a maximum of 700 ACL entries can be set. The maximum is the total number of access ACL entries and default ACL entries.

Migrating to an Advanced ACL type of file system

In the HDI system, a Classic ACL type of file system can automatically be migrated to an Advanced ACL type (which allows file systems to be remounted) or by using the `fsctl` command.

In the HDI system, please note the following when migrating a Classic ACL type of file system, that stores existing shared information, to an Advanced ACL type of file system:

- Migrating by the `XCOPY` command or backup utility

The content displayed in the Classic ACL type of file system might differ from the actual audited access. As `XCOPY` command copies access permissions displayed in ACL properties and the forward audited permission will be changed, the great care should be required when

migrating from a Classic ACL type of file system to an Advanced ACL type of file system.

ACL set by default if there is no inherited ACL

If a file or folder is created on an Advanced ACL type of file system, inheritable ACEs will be searched for from among the ACLs set for the parent folders, and that ACL will be set for the created file or folder. If there is no inheritable ACL in the parent folders, the ACL shown in the following tables will be set by default.

Folder

Table 8-19 Default ACL inherited by a folder

Item	Description
DOS attribute	DOS_ATTR_DIR
ACE inherit flags	None
Owner	Create user
Owner group	Group to which a create user belongs
ACE	<p>If access permissions are not specified for new folders created in a CIFS share or if the full-control access permission is specified for the owner only:</p> <p>Type: Permit Name: Create user's name Access permission: Full control Applies to: This folder only</p> <p>If access permissions are specified for new folders created in a CIFS share:#</p> <p>Type: Permit or deny (depending on the specified mode) Name: Create user, group which a create user belongs to, or other users Access permission: Specified permission Applies to: This folder only</p>

#

For notes that apply to this case, see [Notes on ACLs set by default for new folders and files created in CIFS shares on page 8-40](#).

File

Table 8-20 Default ACL inherited by a file

Item	Description
DOS attribute	Archive
ACE inherit flags	None

Item	Description
Owner	Create user
Owner group	Group to which a create user belongs
ACE	<p>If access permissions are not specified for new files created in a CIFS share or if the full-control access permission is specified for the owner only:</p> <p style="padding-left: 40px;">Type: Permit</p> <p style="padding-left: 40px;">Name: Create user's name</p> <p style="padding-left: 40px;">Access permission: Full control</p> <p>If access permissions are specified for new files created in a CIFS share: #</p> <p style="padding-left: 40px;">Type: Permit or deny (depending on the specified mode)</p> <p style="padding-left: 40px;">Name: Create user, group which a create user belongs to, or other users</p> <p style="padding-left: 40px;">Access permission: Specified permission</p>

#

For notes that apply to this case, see [Notes on ACLs set by default for new folders and files created in CIFS shares on page 8-40](#).

Notes on the case of migrating from Windows

On a Windows system, the default security policy setting grants all users the right of Bypass traverse checking. Therefore, in most cases for an NTFS ACL of Windows, even if the Traverse Folder right is not permitted for a folder's ACL, operation of the files or folders in that folder is possible by specifying the absolute path when you have access permissions for only the files or folders in the folder.

Even as for the file system of HDI, CIFS bypass traverse checking allows you to manipulate a target object (folder or file) via CIFS access by specifying the absolute path if you only have access permissions for that object. You do not need to have access permissions for any higher-level directories up to the target object.

Note that in the file system which carried over from HDI whose versions are earlier than 4.2.0-00, CIFS bypass traverse checking is disabled by default. To manipulate a target object when CIFS bypass traverse checking is disabled, the Traverse Folder permission must be set in the ACLs for all higher-level directories up to the target object.

For details about CIFS bypass traverse checking, see the *Installation and Configuration Guide*.

Changing file attributes

When changing file attributes in an Advanced ACL file system, changes might not be immediately reflected in Explorer. In this case, open the Explorer **View** menu, and then select **Refresh**.

Notes on ACLs set by default for new folders and files created in CIFS shares

The following describes notes on ACLs set by default when access permissions are set for new folders and files created in a CIFS share.

- When the access permissions (rw, ro, or none) set on an HDI system for the new folders and files created in a CIFS share are displayed on CIFS clients, the access permissions for "Allow" might be different between Advanced ACL type file systems and Classic ACL type file systems, depending on the access permission entries (see the following figure).

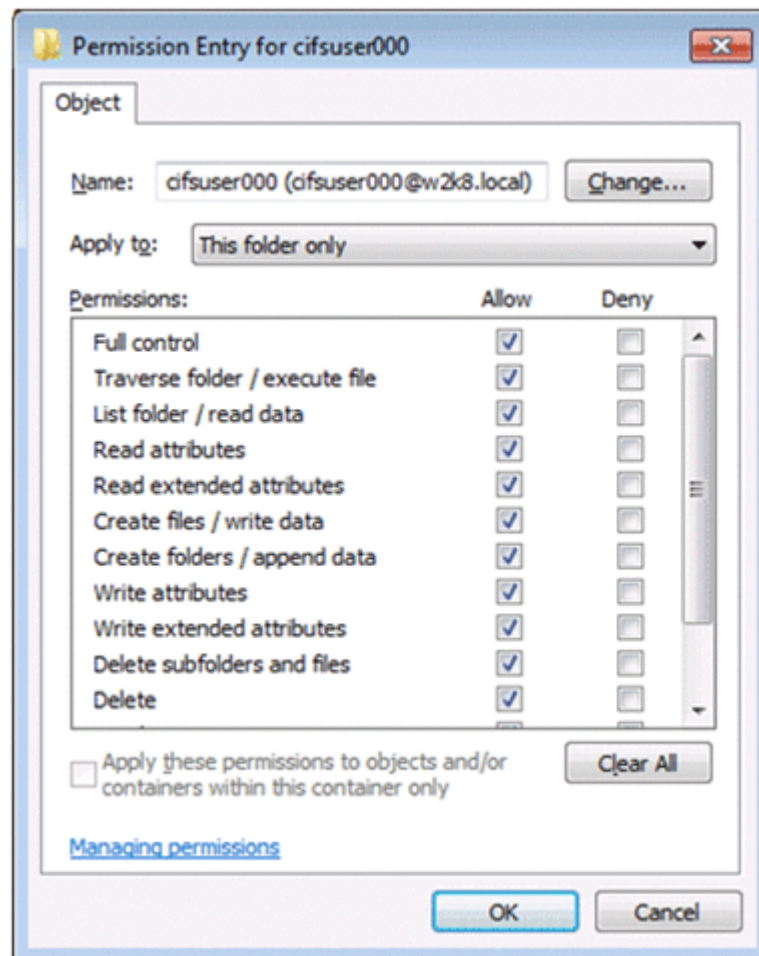


Figure 8-8 Example of access permission entries

The table below shows the differences in the access permissions for "Allow". For entries that are not described in the table, the same access permissions are set for ACLs by default. In the table, "Set" indicates that "Allow" is selected for an access permission entry, and "Not set" indicates that "Allow" is not selected for an access permission entry.

Table 8-21 Displayed access permission entries and specified access permissions (rw, ro, or none) (for folders)

Access permission entries displayed on a CIFS client	Access permissions specified for a CIFS share on a Classic ACL type file system, and NTFS ACL settings	Access permissions specified for a CIFS share on an Advanced ACL type file system, and NTFS ACL settings
Read attributes	rw: Set. ro: Set. none: Not set.	rw: Set. ro: Set. none: Set.
Delete	rw: Set. ro: Not set. none: Not set.	rw: Not set. ro: Not set. none: Not set.
Change permissions	rw: Set. ro: Not set. none: Not set.	If access permissions are being set for Owner: rw: Set. ro: Set. none: Set. If access permissions are being set for Group or Other: rw: Not set. ro: Not set. none: Not set.
Take ownership	rw: Set. ro: Not set. none: Not set.	If access permissions are being set for Owner: rw: Set. ro: Set. none: Set. If access permissions are being set for Group or Other: rw: Not set. ro: Not set. none: Not set.

Table 8-22 Displayed access permission entries and specified access permissions (rw, ro, or none) (for files)

Access permission entries displayed on a CIFS client	Access permissions specified for a CIFS share on a Classic ACL type file system, and NTFS ACL settings	Access permissions specified for a CIFS share on an Advanced ACL type file system, and NTFS ACL settings
Read attributes	rw: Set. ro: Set. none: Not set.	rw: Set. ro: Set. none: Set.
Read permissions	rw: Set. ro: Set. none: Not set.	rw: Set. ro: Set. none: Set.
Change permissions	rw: Not set. ro: Not set. none: Not set.	If access permissions are being set for Owner: rw: Set. ro: Set. none: Set. If access permissions are being set for Group or Other: rw: Not set. ro: Not set. none: Not set.
Take ownership	rw: Not set. ro: Not set. none: Not set.	If access permissions are being set for Owner: rw: Set. ro: Set. none: Set. If access permissions are being set for Group or Other: rw: Not set. ro: Not set. none: Not set.

- For CIFS shares on Advanced ACL type file systems, "Allow" or "Deny" can be set as access permissions. As shown below, "Deny" might be set for some of the access permission entries displayed on CIFS clients, depending on the access permissions set for HDI systems.
 - If "Allow" is not set for "Owner" or "Group" and "Allow" is set for "Other":
 "Deny" is set for "Owner" or "Group".
 - If "Allow" is not set for "Owner" and "Allow" is set for "Group":
 "Deny" is set for "Owner".

Therefore, the ACLs set by default are different between Advanced ACL type file systems and Classic ACL type file systems even if the same access permissions are set for CIFS shares (even if the same combinations of the permissions (rw, ro, or none) are specified for "Owner", "Group", and "Other") on HDI systems. If you do not want "Deny" set for CIFS shares for Advanced ACL type file systems and want the permissions set like the CIFS shares for Classic ACL type file systems, take into account the combinations of access permissions (rw, ro, or none) and configure them according to the following steps. Note that the values in the conversion table are the results of converting the access permission settings into numerical values. They are not the same as the values obtained by expressing `rwX` as an octal number:

- a. Convert the access permissions (rw, ro, and none) to numeric values according to the conversion tables below:

Table 8-23 Conversion table for folders

Access permissions	Targets for which access permissions are set		
	Owner	Group	Other
rw	7	7	7
ro	5	5	5
none	1	1	1

Table 8-24 Conversion table for files

Access permissions	Targets for which access permissions are set		
	Owner	Group	Other
rw	7	6	6
ro	4	4	4
none	0	0	0

- b. Set access permissions so that the relationship between the values specified for the access permissions of Owner, Group, and Other is as follows:

Owner >= Group >= Other
 (>=: greater than or equal to)

The following shows examples of when the relationship of the specified values is satisfied and when it is not satisfied:

When the relationship of the specified values is satisfied:

If the access permissions of a new file created in a CIFS share are as follows:

Owner: `rw`, Group: `ro`, Other: `none`,

the values converted from specified access permissions are as follows:

Owner: 7, Group: 4, Other: 0.

This satisfies the above relationship between the specified values. In this case, the differences of the ACLs set by default are not caused by the differences in ACL types between file systems.

When the relationship of the specified values is not satisfied:

If the access permissions of a new file created in a CIFS share are as follows:

Owner: `rw`, Group: `ro`, Other: `rw`

the values converted from specified access permissions are as follows:

Owner: 7, Group: 4, Other: 6.

This does not satisfy the above relationship between the specified values because the value set for Other is larger than the value set for Group. In this case, "Deny" is set for the access permission entry for writing new files created in the CIFS share. As a result, the owner belonging to the group also cannot update new files.

- Do not set `ro` (converted value: 5 for folders, and 4 for files) or `none` (converted value: 1 for folders, and 0 for files) for Owner as the permission to access the new folders or files created in a CIFS share. If the permissions are set as such, creating files in new folders created in the CIFS share and writing to the new files created in the CIFS share are impossible even if the Owner permission is used.
- If CREATOR OWNER or CREATOR GROUP is displayed when a CIFS share on a file system of the Advanced ACL type and if new folders or files are created in the folder, two types ACEs might be set for the user and the group to which the user belongs as the ACEs of the folders or files regardless whether the access permissions are set for the CIFS share. The conditions are shown below.

If two types of ACEs are set for the same user:

If a CIFS share is configured so that the ACE set for the creator of a new folder or file is inherited and the access permissions or application targets set for the creator's ACE are different with those of CREATOR OWNER, the ACEs for the creator and CREATOR OWNER are set for new folders or files created in the folder.

If two types of ACEs are set for the same group:

If a CIFS share is configured so that the ACE for the group to which the creator of a new folder or file belongs is inherited and the access permissions or application targets set for the group's ACE are different with those of CREATOR GROUP, the ACEs for the group to which the creator belongs and CREATOR GROUP are set for the new folders or files created in the folder.

Adding user ACLs or group ACLs

This subsection describes how to add user or group ACLs.

User or group ACLs can be added by choosing **Add** in the access permission window for a file or folder. The following figure shows the **Select User or Group** window that appears when the **Add** button is clicked.




Figure 8-9 Select User or Group window (Left: When user mapping is not being used; Right: When user mapping is being used)

If you specify ACL settings for a folder created in a CIFS file share by using the file or folder properties window, the location of a user or group you select differs depending on the CIFS service authentication method. The location mentioned above corresponds to the number 1 in the above figure.

When user mapping is not used for local authentication, Active Directory authentication, or NT domain authentication

You must select a user or group for which the host name of an HDI node is displayed in the **From this location** box of the **Select User or Group** window.

Notes:

- When Active Directory authentication or NT domain authentication is selected, a domain controller might be displayed in the **From this location** box. Do not specify it. If it is specified, the ACL setting will not be enabled.
- If you want to display a user (the host name of an HDI node) in the **From this location** box on the left side of [Figure 8-9 Select User or Group window \(Left: When user mapping is not being used; Right: When user mapping is being used\)](#) on page 8-45, register a CIFS user on the File Services Manager **Add User** page of the **Local Users** dialog box.
- To display a group in that box, select the **Apply to CIFS ACL environment** check box when adding a group on the File Services Manager **Add Group** page of the **Local Users** dialog box.
- Depending on the environment, user or group ACLs might not be able to be added in the access permission window for a file or folder. In such cases, use the `dirsetacl` command to specify ACL settings. If you want to add an ACL to a file, use the `dirsetacl` command to specify the ACE inheritance range for the relevant directory so that the required ACEs are inherited by the target file.

When user mapping is used for Active Directory authentication or NT domain authentication

You must select a user or group for which the domain controller is displayed in the **From this location** box in the **Select User or Group** window.

If the domain controller is not displayed for the **From this location** box in the **Select User or Group** window, communication with the domain controller might fail. Possible causes are as follows:

- o An incorrect value is specified for **Domain name (NetBIOS)** in the **Active Directory Authentication** page of the **Access Protocol Configuration** dialog box.
- o On the CIFS client, the IP address of the domain controller cannot be resolved by DNS.
- o The user who is performing CIFS is not a domain user.

The ACL of a local group or local user can be configured by selecting a user or group for which the host name of an HDI node is displayed in the **From this location** box from the **Select User or Group** window.

Notes:

- When you configure the ACL of an HDI local group or local user, use the **Select User or Group** window on the left side of [Figure 8-9 Select User or Group window \(Left: When user mapping is not being used; Right: When user mapping is being used\) on page 8-45](#). A CIFS user must be registered via the **Add User** page of the File Services Manager **Local Users** dialog box to display a user for which the host name of an HDI node is displayed in the **From this location** box.
- To display a group in that box, select the **Apply to CIFS ACL environment** check box when adding a group on the File Services Manager **Add Group** page of the **Local Users** dialog box.
- A CIFS client must be registered in a domain, or no HDI local groups or local users are displayed, even if the host name of an HDI node is displayed in the **From this location** box in the **Select User or Group** window.
- In some environments, you might not be able to add a user or group ACL from the access permission window for a file or folder. In such cases, use the `dirsetacl` command to specify the necessary settings. To add an ACL to a file, use the `dirsetacl` command to set the inheritance range of the access control entries (ACEs) for the relevant directory so that the file inherits the ACEs.

Note that, regardless of the CIFS service authentication format, if you cannot add a user ACL or group ACL from the file or folder access permissions window, you might be able to add the ACL by using one of the following methods:

- Log on to the Windows client as an administrator (a user belonging to the Administrators group).
- Log on to the Windows client as user who is not an administrator (as a user who was individually added to the Windows client).
- When the authentication dialog box for adding a user ACL or a group ACL appears, enter your user name and password for connecting to HDI system.

- Make sure the user name and password you use to connect to HDI system and those you use for logging on to the Windows client are the same.
 - Connect to HDI system by specifying the host name from the Windows client.
 - Make sure the Windows client is a participant in the Active Directory domain.
- If you do this, no problem will occur even if the Windows client has no relationship to the HDI system authentication method or to the domain in which HDI system participates.

File attributes

This section describes operations for file attributes of a CIFS share from a CIFS client.

Setting and checking file attributes from a CIFS client

This subsection describes how to, from a CIFS client, set and view the attributes of shared files.

Users who can set file attributes

In the HDI system, only CIFS administrators registered in File Services Manager and those who have Write permission for a file or directory can set file attributes for a file. The owner of a file who does not have Write permission on the file cannot set file attributes for the file.

Whether file attributes can be set

Whether file attributes that have been set from a CIFS client can be set in an HDI system depends on the ACL type, as shown in the following table.

Table 8-25 Whether file attributes can be set in an HDI system

File attribute	Description	Applies to Classic ACL	Applies to Advanced ACL
Read Only	This attribute indicates that the file is non-writable or cannot be moved. This attribute may also be called the <i>write-protected</i> attribute.	Available	Available
System	This attribute indicates an important file required to run the system. Normally, the files that have this attribute must not be moved or changed.	Unavailable	Available
Hidden	Normally, the files that have this attribute are invisible from the shell. However, Explorer can be set to make these files visible.	Unavailable	Available

File attribute	Description	Applies to Classic ACL	Applies to Advanced ACL
Archive	This attribute indicates that the file has been updated after it was last backed up.	Unavailable	Available [#]
Compressed	This attribute is available only when the file system is NTFS. This attribute indicates a file that has been compressed at the file system level.	Unavailable	Unavailable
Encrypted	This attribute is available only when the file system is NTFS. The files that have this attribute are encrypted for high security. This attribute cannot be set together with the Compressed attribute. It is not recommended that you back up the files that have this attribute by using a backup tool that does not support this attribute. If you do it, the files are decrypted when they are backed up.	Unavailable	Unavailable
Directory	This attribute indicates whether the file is a directory or an ordinary file.	Available	Available
Offline	This attribute indicates that the file is a stub file.	Available	Available
Normal	This attribute indicates that the file is an ordinary file with no attributes set.	Available	Available
Temporary	This attribute indicates that the file is a temporary file.	Unavailable	Unavailable
SparseFile	This attribute indicates that the file is a sparse file.	Unavailable	Unavailable
ReparsePoint	This attribute indicates that the file contains a reparse point.	Unavailable	Unavailable
NotContentIndexed	This attribute indicates that the index service is not enabled for the file.	Unavailable	Unavailable

#

For details, see [Notes on the archive attribute on page 8-49](#).

Notes on sharing a file or directory with NFS

The following are notes on sharing a file or directory with the CIFS and NFS services:

- From an NFS client, if write permissions are deleted for the owner, a group, or other users of a file or directory, the attribute settings are set as read-only for CIFS clients.
- Be aware that if the read-only attribute is set for a file or directory from a CIFS client, the setting does not take effect for NFS clients.

- If the name of a file or directory created by an NFS client begins with a period (.), the hidden file attribute is added in the CIFS share.

Notes on the archive attribute

On an Advanced ACL type of file system, the archive attribute will not set on general files or symbolic files by changing their names or moving them.

Notes on the read-only attribute

Even a CIFS administrator, who is specified in **CIFS administrator name(s)** in the **CIFS Service Management** page (**Setting Type:** *Administration*) of the **Access Protocol Configuration** dialog box, is not able to use Windows API to delete a file or folder that has the read-only attribute.

Offline attribute

In an HDI system, if a file is migrated to an HCP system or is on-demand imported from another file server and then turned into a stub file, the file is managed as a file with the offline attribute. For details about on-demand importing, see the *Installation and Configuration Guide*. Note that the offline attribute cannot be set from CIFS clients.

When files with the offline attribute are displayed in Explorer on a CIFS client, a mark will be present in the lower-left corner of the icon, such as the **x** mark in Windows Server 2012 R2. In the Explorer Attributes column, the offline attribute is designated by an **o**. This icon with the mark might not be displayed for shortcut files with the `offline` attribute. Use the **Type** column in Explorer to determine whether a file is a shortcut. When a list of files is displayed using the command prompt, the size of a file that has the offline attribute appears in parentheses.

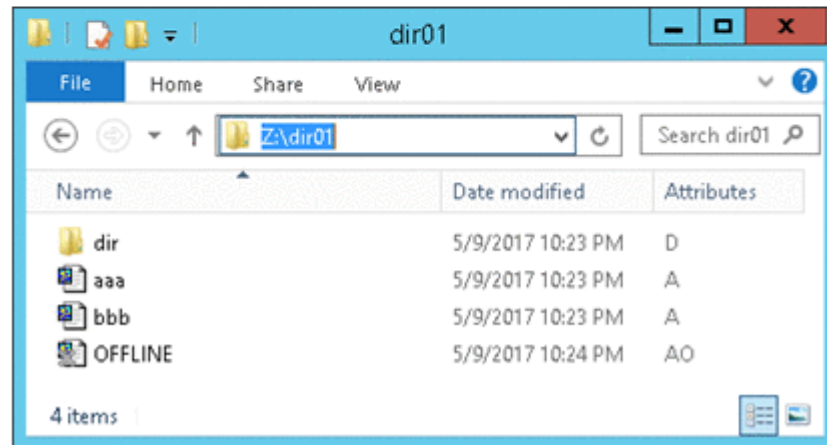
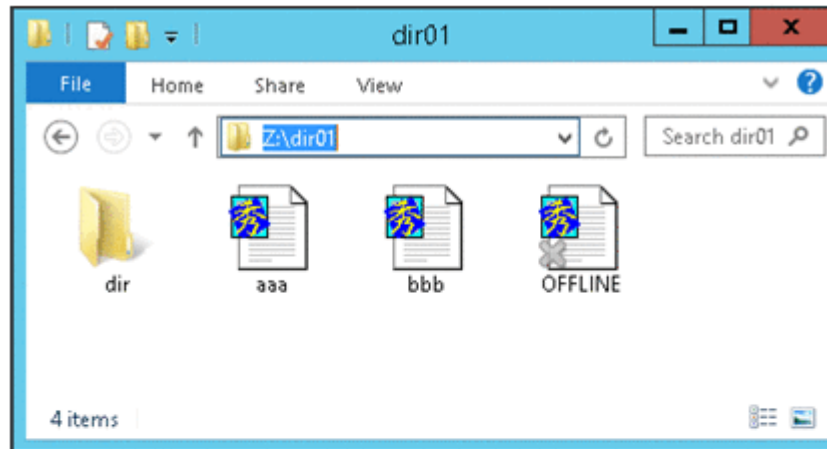


Figure 8-10 Offline attribute as displayed in Explorer




Figure 8-11 Offline attribute as displayed from the command prompt

Extended attributes in Windows

In Windows, extended attributes are used to manage the items displayed in the **General** tab of the Properties window. However, some extended

attributes cannot be used for migration to an HDI system. This is because, depending on the application, extended attributes are stored in an NTFS-specific area called a named stream, and the stream data cannot be copied to file systems other than an NTFS volume.

The following table summarizes the storage locations of extended attributes for major applications.

Table 8-26 Storage locations of extended attributes

Application	Storage locations	Usable in HDI
Microsoft Word	Main data stream	Yes
Microsoft Excel	Main data stream	Yes
Microsoft PowerPoint	Main data stream	Yes
Notepad	SummaryInformation data stream [#]	No
WordPad	SummaryInformation data stream [#]	No
Zip archiver	SummaryInformation data stream [#]	No

#

In the NTFS file system, one file consists of multiple data streams. Of these data streams, the actual file data is stored in the data stream called the main data stream (or anonymous data stream). In a file system other than NTFS, only the main data streams can be accessed. The data streams other than the main data stream are called named data streams. The SummaryInformation data stream is one of the named data streams.

As described above, HDI systems do not support the setting, retrieval, and migration of extended attributes. However, HDI systems support the setting of access permissions related to extended attributes and the checking of access permissions.

Timestamps

This section describes the timestamps for those files accessed in a CIFS file share.

File access date and time

Whether the file access date and time is to be updated can be set using the **Mount File System** dialog box of File Services Manager. For details about how to specify this setting, see the *Administrator's Guide*. The access date and time are updated even when the file properties window is opened.

File modified date and time

This subsection provides notes on the dates and times associated with files being modified.

- When a CIFS client has moved a folder within a CIFS file share, the folder modified date and time are changed to the time at which the CIFS client performed the folder move operation.
- Depending on the specifications of an application such as Microsoft Excel, the access time of a file might be updated when the file is updated even though **Record last access time** in the **Mount File System** dialog box is not set to **Yes**.

File creation date and time

The creation date and time of a file may be updated when the file is updated or when file attributes such as the size or permissions are changed. This is because when an HDI system is set not to hold file creation dates and times, that HDI system returns the update date/time, access date/time, or attribute change date/time, whichever is the earliest, to the CIFS client as the file creation date and time.

File timestamp resolution

This subsection describes the ways in which file timestamp are managed.

File timestamp management method

The following table shows the differences in file timestamp management between the HDI system and Windows (NTFS).

Table 8-27 File timestamp management method

Item	Windows	HDI system
Base time	1601 (year)	1970 (year)
Storage size	8 bytes	4 bytes ^{#1}
Resolution	100 nanoseconds	100 nanoseconds ^{#2}

#1

In a WORM file system, the storage size of a file's access date and time information is 8 bytes.

#2

In a WORM file system, a file's access date and time is measured to the second.

File timestamp update resolution

The following table shows the differences in the file timestamp update resolution between the HDI system and Windows.

Table 8-28 File timestamp update resolution

Timestamp type	Windows	HDI system
File access time	1 hour	100 nanoseconds
File update time	100 nanoseconds	100 nanoseconds
File creation time	100 nanoseconds	100 nanoseconds

Note for granting file timestamp update permission

Even a user, who is specified in **File timestamp changeable users** in the **CIFS Service Management** page (**Setting Type:** *Security*) of the **Access Protocol Configuration** dialog box, is not able to update the timestamp of a file that has the read-only attribute.

Displaying disk capacity

If you link to the HCP system at the CIFS share level, the hard quota of the capacity for the namespace is displayed as the CIFS client disk capacity, not the capacity of the file system capacity. When the hard quota setting for the namespace is changed, the CIFS client's disk capacity is updated during migrations. Note that if the settings of the hard namespace quotas cannot be obtained or 0 is set as the hard quota, the block capacity and the number of inodes of the file system is displayed as the disk capacity of the CIFS client.

When the capacity of a share is not limited based on the hard quota for the namespace, CIFS users can view the disk capacity of a share in the Properties dialog of the share. If quotas are set for a file system or directory that contains the share, the value of the disk capacity shows the amount limited by the quota settings. However, when a CIFS administrator registered in File Services Manager views the capacity of a file share, the actual capacity of the share is displayed, not the capacity that is based on any file system or directory quotas that might exist for the share.

The following subsections explain the relationship between quota settings and the available disk capacity.

The quota management functionality can set the limit of blocks number or inodes number that users can use. The shared size will be affected in a display when using CIFS share.

Quota managing in the HDI system differs from Windows' as listed below:

Security information available on quota management functionality

Security information here means file's owner or owner group. On the HDI system, the quota management functionality is applied to the file owner and owner group so that files cannot be used when either (or both) of them exceeds their limit. On the other hand, Windows applies to only owner for quota functionality.

Types available on default quota management functionality

Default quota management functionality is for automatically applying quota functionality to a user who does not have a limit on the number of blocks or inodes they can use. On the HDI system, the default quota functionality is applied to the owner, only when the owner is a user (and not when owner is a group). However, in Windows, the default quota functionality is always applied to the owner (whether they are a user or a group).

The following table shows the differences between the HDI system and Windows.

Table 8-29 Differences between the HDI system and Windows on quota functionality

Security information		Preventing or watching		Default quota functionality	
		HDI	Windows	HDI	Windows
Owner	User	Available	Available	Available	Available
	Group	Available	Available	Not available	Available
Owner group		Available	Not available	Not available	Not available

The following table shows the differences in security information between the HDI system and Windows. The following table shows the size of CIFS shares (as shown in the properties window) and whether quota functionality is applied in the cases of when user quota is applied to the owner, or when group quota is applied to the owner's group, or when directory quota (subtree directory quota in HDI systems) is applied to the top of shared directories.

Table 8-30 Differences between the HDI system and Windows on security information

Security information	Applies quota functionality		Shared size	
	HDI	Windows	HDI	Windows
User quota is applied to the owner	Yes	Yes	Quota upper limit	Quota upper limit
Group quota is applied to the owner's group	Yes	No	Quota upper limit	Total capacity of file system
Directory quota is applied to the top of shared directories	Yes	Yes	Quota upper limit	Quota upper limit

In an HDI system, when quotas are set for file systems, disk space is based on the user quotas, default quotas, and group quotas. When quotas are set for directories, disk space is based on the subtree user quotas, subtree default quotas, subtree group quotas, and subtree directory quotas.

Advanced ACL file systems do not support quotas or default quotas for groups that are owners. A group quota is evaluated from the total of the capacity of

the files whose owners are a specific group and of the capacity of the files whose file owner groups are the specific group.

Even if the disk space viewed from the CIFS client indicates that there is sufficient free space to write data, an error might occur due to insufficient disk space. In this case, use the `fslist` command to check the available space in the file system block. In addition, use the `quotaget` and `stquota` commands to check the block usage and the inode usage for the user or affiliated group.

Whether the quotas can be checked on a CIFS client

You can check the quotas set in the HDI system from a CIFS client by viewing the disk capacity. The following table shows whether the quota settings can be checked from a CIFS client.

Table 8-31 Whether the quotas set in the HDI system can be checked from a CIFS client

Quota settings			Checkable from a CIFS client
Subtree user quota	Block capacity	Soft limit	Yes
		Hard limit	Yes
		Grace period	No
	inode	Soft limit	No
		Hard limit	No
		Grace period	No
Subtree default quota	Block capacity	Soft limit	Yes
		Hard limit	Yes
	inode	Soft limit	No
		Hard limit	No
Subtree group quota	Block capacity	Soft limit	Yes
		Hard limit	Yes
		Grace period	No
	inode	Soft limit	No
		Hard limit	No
		Grace period	No
Subtree directory quota	Block capacity	Soft limit	Yes
		Hard limit	Yes
		Grace period	No
	inode	Soft limit	No
		Hard limit	No
		Grace period	No

Quota settings			Checkable from a CIFS client
User quota	Block capacity	Soft limit	Yes
		Hard limit	Yes
		Grace period	No
	inode	Soft limit	No
		Hard limit	No
		Grace period	No
Default quota	Block capacity	Soft limit	Yes
		Hard limit	Yes
	inode	Soft limit	No
		Hard limit	No
Group quota	Block capacity	Soft limit	Yes
		Hard limit	Yes
		Grace period	No
	inode	Soft limit	No
		Hard limit	No
		Grace period	No
Namespace quota	Block capacity	Hard limit	Yes

As shown in the above table, only the quota settings related to block capacity can be displayed on a CIFS client. Note that whether the displayed disk capacity is the soft limit or the hard limit cannot be checked.

The quota value displayed for **Disk capacity** is applicable to the directory to which a drive is allocated. Note that the value displayed for **Disk capacity** varies depending on the quota settings and disk usage. For details, see [Disk capacity displayed in accordance with disk usage on page 8-57](#) and [Disk capacity displayed when multiple quotas are set on page 8-59](#).

Also note that an incorrect value might be displayed for the disk capacity if separate drives are allocated for multiple directories. Use the `cifsoptset` command to set 0 for `dfree_cache_time` so that the information about the free disk capacity is not cached unless there is any impact, such as a response delay, on the client. For details about how to use the `cifsoptset` command to prevent the disk capacity information from being cached, see the *CLI Administrator's Guide*.

The following figures are examples of disk capacity information displayed on a CIFS client. The information enclosed in the frame varies depending on the quota settings and disk usage.




Figure 8-12 Disk capacity displayed when no quota is set

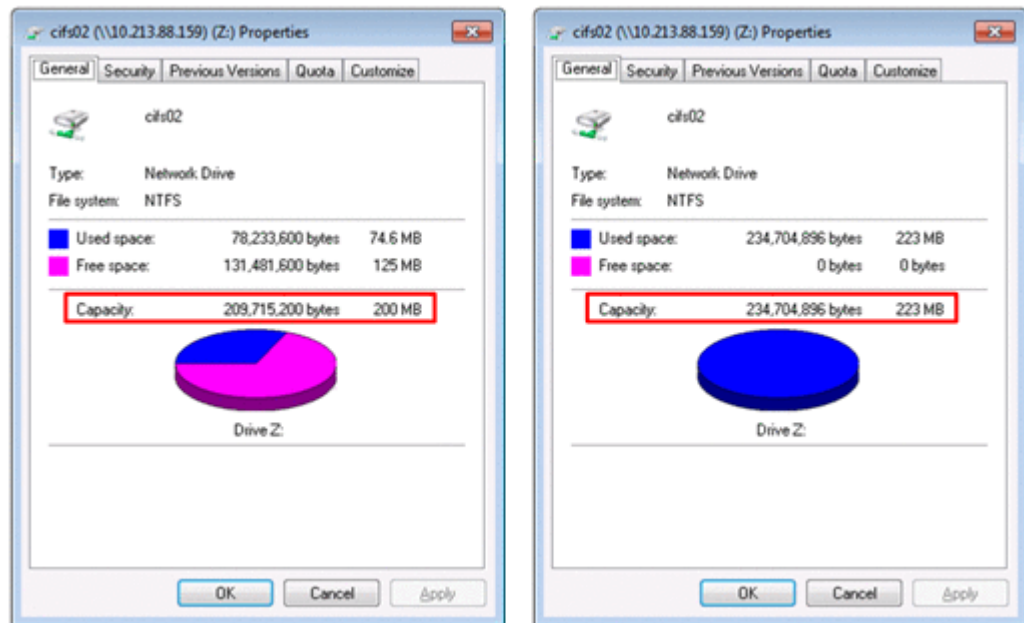


Figure 8-13 Disk capacity displayed when a quota is set (Left: quota not exceeded, right: quota exceeded)

Disk capacity displayed in accordance with disk usage

This subsection describes how disk capacity is displayed in a CIFS client when the block capacity or the number of inodes is set as a quota in an HDI system. When the capacity of a share is limited, based on the hard quota for the namespace, the block size that was specified as the migration destination's hard quota of the capacity for the namespace is displayed as the disk capacity.

When a block capacity quota is set

When a block capacity quota is set in an HDI system, the disk capacity is displayed on a CIFS client, based on the disk usage as shown in the following table.

Table 8-32 Disk capacity displayed on a CIFS client when a block capacity quota is set in an HDI system

Quota value		Usage	Disk capacity
Soft limit	Hard limit		
Not set	Not set	--	File system capacity
Not set	Set	At or above the hard limit	Block usage
		Below the hard limit	Hard limit on the block capacity
Set	Not set	At or above the soft limit	Block usage
		Below the soft limit	Soft limit on the block capacity
Set	Set	At or above the hard limit	Block usage
		At or above the soft limit	Block usage
		Below the hard limit	
		Below the soft limit	Soft limit on the block capacity

Legend:

--: Not applicable

When a quota for number of inodes is set

When a quota for the number of inodes is set in an HDI system, the disk capacity is displayed on a CIFS client, based on the disk usage as shown in the following table. Note, however, that the number of inodes that has been set cannot be checked from the client.

Table 8-33 Disk capacity displayed on a CIFS client when a quota for the number of inodes is set in an HDI system

Quota value		Usage	Disk capacity
Soft limit	Hard limit		
Not set	Not set	--	File system capacity
Not set	Set	At or above the hard limit	Block usage
		Below the hard limit	File system capacity

Quota value		Usage	Disk capacity
Soft limit	Hard limit		
Set	Not set	At or above the soft limit	Block usage
		Below the soft limit	File system capacity
Set	Set	At or above the hard limit	Block usage
		At or above the soft limit	Block usage
		Below the hard limit	Block usage
		Below the soft limit	File system capacity

Legend:

--: Not applicable

Disk capacity displayed when multiple quotas are set

This subsection explains the disk capacity values displayed on a CIFS client to which multiple quotas apply.

The HDI system

In an HDI system, the disk capacity displayed on a CIFS client depends on whether the disk usage quota has been reached.

When the amount of used disk space has not reached a quota limit

If no quotas for block capacity and number of inodes used that apply to a CIFS client have not been reached, disk capacity is displayed based on the rules shown in the following table.

Table 8-34 Disk capacity displayed when multiple quotas are set (when no quotas have been reached)

Name space quota	Quotas set for each directory (subtree quotas)				Quotas set for each file system			Disk capacity
	Subtree user quota	Subtree default quota	Subtree group quota	Subtree directory quota	User quota	Default quota	Group quota	
Block capacity limited	--							Block capacity limit value for the namespace quota

Name space quota	Quotas set for each directory (subtree quotas)				Quotas set for each file system			Disk capacity
	Subtree user quota	Subtree default quota	Subtree group quota	Subtree directory quota	User quota	Default quota	Group quota	
Block capacity not limited	Block capacity limited	--						Block capacity limit value for the subtree user quota
Block capacity not limited	No quota set	Block capacity limited	--					Block capacity limit value for the subtree default quota
Block capacity not limited			Block capacity limited	--				Block capacity limit value for the subtree group quota
Block capacity not limited				Block capacity limited	--			Block capacity limit value for the subtree directory quota
Block capacity not limited					Block capacity limited	--		Block capacity limit value for the user quota
Block capacity not limited					No quota set	Block capacity limited	--	Block capacity limit value for the default quota
Block capacity not limited						Block capacity limited		Block capacity limit value for the group quota
Block capacity not limited								File system size

Legend:

"--" indicates that the displayed disk capacity does not depend on whether the quota is set.

Notes:

- "Limit value" indicates the soft limit when the soft limit is set, and indicates the hard limit in other cases.

- "Block capacity not limited" means that the capacity of the share is not limited based on the hard quota for the namespace.

"Limit value" indicates the soft limit when the soft limit is set, and indicates the hard limit in other cases.

When a disk usage quota has been reached

If any quota for block capacity or number of inodes used that apply to a CIFS client has been reached, disk capacity is displayed based on the rules shown in the following table.

Table 8-35 Disk capacity displayed when multiple quotas are set (when any quota is reached)

Name space quota	Quotas set for each directory (subtree quotas)				Quotas set for each file system			Disk capacity
	Subtree user quota	Subtree default quota	Subtree group quota	Subtree directory quota	User quota	Default quota	Group quota	
Usage exceeds the limit	--							Block capacity limit value for the namespace quota
Usage is within the limit	--							Block capacity limit value for the namespace quota
Block capacity not limited	Usage exceeds the limit	--						Block usage for the subtree user quota
Block capacity not limited	Usage is within the limit	Usage exceeds at least one limit						Block capacity limit value for the subtree user quota
Block capacity not limited				Usage exceeds the limit	--			Block usage for the subtree group quota
Block capacity not limited				Usage is within the limit	Usage exceeds at least one limit			Block capacity limit value for the subtree group quota

Name space quota	Quotas set for each directory (subtree quotas)				Quotas set for each file system			Disk capacity
	Subtree user quota	Subtree default quota	Subtree group quota	Subtree directory quota	User quota	Default quota	Group quota	
Block capacity not limited				Usage exceeds the limit	--			Block usage for the subtree directory quota
Block capacity not limited				Usage is within the limit	Usage exceeds at least one limit			Block capacity limit value for the subtree directory quota
Block capacity not limited				Usage exceeds the limit	--			Block usage for the user quota
Block capacity not limited				Usage is within the limit	Usage exceeds at least one limit			Block capacity limit value for the user quota
Block capacity not limited						Usage exceeds the limit	Block usage for the group quota	
Block capacity not limited						Usage is within the limit	Block capacity limit value for the group quota	

Legend:

"--" indicates that the displayed disk capacity does not depend on whether the quota is set.

Notes:

- "Limit value" indicates the soft limit when the soft limit is set, and indicates the hard limit in other cases.
- "Block capacity not limited" means that the capacity of the share is not limited based on the hard quota for the namespace.

Windows server

When a block capacity quota is set on a Windows server, the disk capacity is displayed on a CIFS client as shown in the following table.

Table 8-36 Disk capacity displayed on a CIFS client where the file server is Windows-based and disk quotas are set on the file server

Directory quota		Disk quota		Usage	Disk capacity		
Soft limit set	Hard limit set	[Warning level] set	[Limit disk space] set				
No	No	No	No	--	Volume capacity		
		Yes	Yes	At or above the value set for [Limit disk space]	Value set for [Limit disk space]		
				Below the value set for [Limit disk space]	Value set for [Limit disk space] [#]		
				Below the value set for [Warning level]	Value set for [Limit disk space] [#]		
No	Yes	No	No	At or above the value set for [Hard limit]	Value set for [Hard limit] [#]		
				Below the value set for [Hard limit]	Value set for [Hard limit]		
		Yes	Yes	At or above the value set for [Limit disk space]	The smaller of the [Limit disk space] and [Hard limit] values [#]		
				At or above the value set for [Hard limit]	The smaller of the [Limit disk space] and [Hard limit] values [#]		
				At or above the value set for [Warning level]	The smaller of the [Limit disk space] and [Hard limit] values [#]		
				Below the value set for [Warning level]	The smaller of the [Limit disk space] and [Hard limit] values [#]		
		Yes	No	No	No	At or above the value set for [Soft limit]	Value set for [Soft limit] [#]
						Below the value set for [Soft limit]	Value set for [Soft limit]
Yes	Yes			At or above the value set for [Limit disk space]	The smaller of the [Limit disk space] and [Soft limit] values [#]		

Directory quota		Disk quota		Usage	Disk capacity
Soft limit set	Hard limit set	[Warning level] set	[Limit disk space] set		
				At or above the value set for [Soft limit] Below the value set for [Limit disk space]	The smaller of the [Limit disk space] and [Soft limit] values [#]
				At or above the value set for [Warning level] Below the value set for [Soft limit]	The smaller of the [Limit disk space] and [Soft limit] values [#]
				Below the value set for [Warning level]	The smaller of the [Limit disk space] and [Soft limit] values [#]
Yes	No	No	No	At or above the value set for [Soft limit]	Value set for [Soft limit] [#]
				Below the value set for [Soft limit]	Value set for [Soft limit]
		Yes	Yes	At or above the value set for [Limit disk space]	The smaller of the [Limit disk space] and [Soft limit] values [#]
				At or above the value set for [Soft limit] Below the value set for [Limit disk space]	The smaller of the [Limit disk space] and [Soft limit] values [#]
				At or above the value set for [Warning level] Below the value set for [Soft limit]	The smaller of the [Limit disk space] and [Soft limit] values [#]
				Below the value set for [Warning level]	The smaller of the [Limit disk space] and [Soft limit] values [#]

Legend:

--: Usage does not depend on whether a quota setting is specified.

#

A Windows file server always pass the value set for **Limit disk space to field**, the value varies in accordance with disk usage in an HDI system.

WORM files

WORM is a function that makes files within a specific file system (called a *WORM file system*) read-only and prevents data from being changed or

deleted either for a fixed time period or indefinitely. A file that has been set to a WORM state is called a *WORM file*.

A WORM file has the characteristics below. For details on WORM file systems, see the *Installation and Configuration Guide*.

To set or extend a retention period for a file, use a custom application that you have created. For details on the APIs used for creating custom applications, see [Appendix F, APIs for WORM operation on page F-1](#).

- The file cannot be written to.
No one can write to a WORM file even if the Write access permission in the file's ACL is set to `Allow`.
You can assign or remove the read-only attribute for WORM files.
- Whether a file is WORM file depends on the file attribute setting, not on the ACL.
A file becomes a WORM file when Read-only is set as the file attribute. If Read-only is permitted only in a file's ACL, the file is not changed to a WORM file.
- Finite retention periods and infinite retention periods for WORM files.
A retention period (the period of time for which a file is retained) is set to the `atime` attribute for WORM files.
A retention period in which data cannot be modified or deleted for a set period of time is called a finite retention period. A time in the future must be set for a finite retention period. For WORM files with a finite retention period, the `atime` attribute is always a time in the future from when the finite retention period.
A retention period in which data cannot be modified or deleted for an indefinite amount of time is called an infinite retention period. A time that is 24 hours or more before the current time must be set for an infinite retention period. For WORM files with an infinite retention period, the `atime` attribute is always 24 hours or more before the time when the infinite retention period was set.
- A retention period set for a WORM file can be extended.
If a retention period is finite, you can extend the set retention period, but you cannot shorten the retention period.
If a retention period is infinite, you cannot modify the set retention period. In addition, you cannot change a finite retention period to an infinite retention period.
- `atime` is measured in seconds.
In a WORM file system, `atime` is measured in seconds.
- A WORM file's `atime` is no longer updated.
When a CIFS client accesses a non-WORM file, `atime` is updated. However, when a CIFS client accesses a WORM file for which a retention period has been set, `atime` is not updated.
- To delete a WORM file, the read-only attribute must be removed.

If a set retention period for a WORM file has expired, the file can be deleted by removing the read-only attribute. However, the data cannot be modified.

Access Control by using ABE

Access based enumeration (ABE) is a function that shows or hides a file or folder depending on whether **Read** permission is set when a CIFS client attempts to display a list of files and folders. If ABE is enabled, the files and folders for which **Read** permission is not set are not displayed on the CIFS client. Note that, if you enable ABE, listing of files and folders on the CIFS client might be slower.

How ABE controls whether to display files and folders

This subsection uses examples to explain how ABE controls whether to display files and folders. For details about how to specify the ABE settings, see the *Administrator's Guide*.

The following table describes the files and folders displayed and not displayed on the client according to whether **Read** permission is set and whether ABE is enabled.

Table 8-37 Files and folders displayed and not displayed on the client according to whether Read permission is set and whether ABE is enabled

Folder or file name	Read permission setting for the folder or file	Whether the folder or file is displayed on the client	
		When ABE is enabled	When ABE is disabled
dir1	Allow	Displayed	Displayed
file11	Allow	Displayed	Displayed
file12	Deny	Not displayed	Displayed
dir2	Deny	Not displayed	Displayed
file21	Allow	Not displayed	Not displayed [#]
file22	Deny	Not displayed	Not displayed [#]

#

This file is in the `dir2` folder, for which **Read** permission is not set. The file is not displayed because a list of files under the folder cannot be acquired.

The following figures show examples of displaying folders and files on a CIFS client. If ABE is enabled, the `file12` file and the `dir2` folder are not displayed because **Read** permission is not set for them ([Figure 8-14 Example when ABE is enabled on page 8-67](#)). If ABE is disabled, the `dir1` and `dir2` folders and the `file11` and `file12` files contained in the `dir1` folder are displayed

regardless of the access permission setting. However, the files in the `dir2` folder are not displayed because **Read** permission is not set for the folder (Figure 8-15 Example when ABE is disabled on page 8-67). Note that ABE controls only whether to display files and folders. When a file is not displayed but its path is known, the file can be accessed if an access permission for the file is granted. For example, if **Traverse folder** permission is set by the ACL for the `dir2` folder in Table 8-37 Files and folders displayed and not displayed on the client according to whether Read permission is set and whether ABE is enabled on page 8-66, the `file21` file can be accessed by specifying its path.




Figure 8-14 Example when ABE is enabled

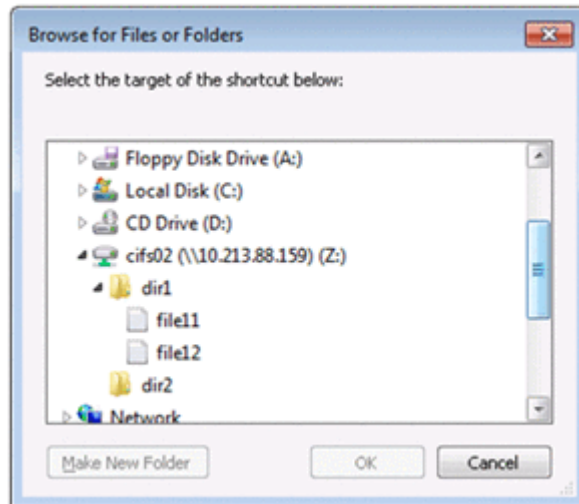


Figure 8-15 Example when ABE is disabled

Shortcut files are also subject to control by ABE. Therefore, if ABE is enabled, shortcut files for which **Read** permission is not set are not displayed. Even when a shortcut file is displayed, if **Read** permission is not set for the entity file, the entity file is not displayed.

Note that ABE has no effect on the CIFS administrator registered in File Services Manager because the CIFS administrator is a root user.

If ABE is disabled, folders and files for which no access permissions are set are also displayed. If an attempt is made to access a folder or file for which no access permissions are set, the attempt is rejected, and a screen such as the following appears.




Figure 8-16 Example of the result of an attempt to access a folder or file for which no access permissions are set

About Read permission required for displaying files and folders when ABE is enabled

For a file or folder to be displayed when ABE is enabled, **Allow** must be selected for **Read** permission in the ACL settings for the file or folder. However, this permission additionally consists of the following five permissions, and the file or folder is not displayed if at least one of these permissions is missing:

- **List folder/read data**
- **Read attributes**
- **Read extended attributes**
- **Read permissions**
- **Synchronize[#]**

#

Synchronize permissions are automatically set when files and folders are created, and cannot be set by using the GUI.

Whether **Read** permission required for displaying a file or folder is granted is determined by the logical sum of the above five permissions. For example, if a user who is granted only **List folder/read data** permission belongs to a group that is granted **Read attributes**, **Read extended attributes**, **Read permissions** and **Synchronize** permissions, the file or folder is visible to the user.

Note that the owner of a file has the privilege of manipulating the access permission settings of the file. Therefore, the owner of a file or folder can view and operate it without **Read permissions** permission if the other four permissions are set.

Also note that the detailed ACL settings shown above cannot be specified if the HDI file system uses Classic ACL because it is a POSIX-compliant ACL.

The following figure shows the windows that display permissions required for displaying a file or folder when ABE is enabled.




Figure 8-17 Windows that display permissions required for displaying a file or folder when ABE is enabled (left: basic permission settings, right: advanced permission settings)

Restrictions on files and folders on CIFS shares

This section describes restrictions on files and folders on CIFS shares.

A VHD (virtual hard disk) file or VHDX file that is used as a backup image or as a Hyper-V virtual hard disk in Windows can neither be created nor used on a CIFS share in HDI.

MMC Linkage

You can manage a CIFS share from Microsoft Management Console (MMC) by using the Shared Folders function in the **Computer Management** window, which is a Windows administrative tool. After adding the Shared Folders snap-in to MMC, CIFS administrators can manage CIFS shares in the HDI system and connections from CIFS clients to CIFS shares, and files on the CIFS shares opened by CIFS clients.

This chapter provides notes on CIFS share management from MMC.

- [Linking an HDI system with MMC](#)
- [Operations required to link with MMC \(for system administrators\)](#)
- [Linking to MMC \(for CIFS administrators\)](#)
- [Before using the administrative share](#)
- [CIFS share management from MMC](#)
- [Session management from MMC](#)
- [Managing open files from MMC](#)
- [Share-level ACLs](#)
- [Notes on using MMC](#)

Linking an HDI system with MMC

An HDI system can be linked with following MMC versions:

- MMC 1.2
- MMC 2.0
- MMC 3.0

The HDI system enables you to use the Windows shared folder features listed in the following table.

Table 9-1 List of Windows shared folder features

Shared Folders function	
Shares	Creating CIFS shares
	Deleting CIFS shares
	Changing the comment for a CIFS share
	Restricting the number of CIFS share users
	Specifying whether CIFS share caching is enabled
	Displaying a list of CIFS shares ^{#1 #2}
	Setting up the share level ACLs
	Setting up the ACLs for shared folders ^{#1}
Sessions	Displaying a list of sessions ^{#1 #2}
	Disconnecting sessions
Open Files	Displaying a list of files that are currently open ^{#1 #2}
	Closing files

#1:

In an HDI system, in addition to CIFS administrators, end users can also use this function.

#2:

In an HDI system, some information is not displayed correctly. For details, see [Table 9-2 Items that can be viewed in the CIFS share list and their reliability in an HDI system on page 9-4](#), [Table 9-5 Items that can be viewed in the list of sessions and their availability in an HDI system on page 9-8](#) and [Table 9-6 Items displayed in the list of open files in an HDI system and their availability on page 9-9](#).

Operations required to link with MMC (for system administrators)

To link an HDI system to MMC, the system administrator must perform the following operations in File Services Manager beforehand. For details on each operation, see the *Administrator's Guide*:

- Create and mount file systems.
The system administrator must create the file systems to be managed by MMC and then mount them.
- Confirm that the CIFS service is running.
In the **List of Services** page, the system administrator must confirm that the CIFS service is running.
- Check the settings of the CIFS administrator.
In the **CIFS Service Management** page (**Setting Type: Administration**), make sure that the CIFS administrator is specified.

Linking to MMC (for CIFS administrators)

To link an HDI system to MMC, the CIFS administrator must connect to the HDI system by adding the Shared Folders snap-in to MMC.

The procedure in this subsection is written for Windows 7.

The CIFS administrator must use either of the following methods to open MMC:

- Click **Start, Accessory, Run**, type `mmc`, and then click the **OK** button.
- Enter `mmc` at the command prompt, and then press the **Enter** key.

To add the Shared Folders snap-in to MMC, and then connect to the HDI system:

1. On the main toolbar, click **File**, and then **Add/Remove Snap-in**.
2. In the **Add or Remove Snap-ins** dialog box, select the **Shared Folders** snap-in from the **Available Snap-ins** list, and then click the **Add** button.
3. In the **Shared Folders** dialog box, select the **Another Computer** radio button, specify the virtual IP address or host name of a node, and then click the **Finish** button.
4. In the **Add or Remove Snap-ins** dialog box, click the **OK** button.
For the node you selected in step 3, the Shared Folders snap-in is inserted into the console tree.
5. Save the console.
After saving the console, you can use it by selecting **Start, All Programs**, and then **Administrative Tools**.

Before using the administrative share

When using MMC to manage CIFS shares in the HDI system, the CIFS administrator can use the following administrative share (the default share) to access the CIFS shares. Note that the administrative share can also be used when a CIFS client views the CIFS shares.

- Share name: `C$`
- Share path: `/mnt`

Note the following when using the administrative share:

- You cannot use the administrative share from File Services Manager.
- You cannot create, delete, or update a file system immediately under the administrative share.
- If a file system immediately under the administrative share is being unmounted or is blocked, you cannot display a list of folders belonging to the file system or create, delete, or update such folders.
- You cannot create, delete, or manage CIFS shares for a file system that belongs to a resource group on the other node.

CIFS share management from MMC

The CIFS administrator can use MMC to create, delete, or update a CIFS share for the connected file system.

When managing CIFS shares in an HDI system by using MMC, invalid information might be displayed for some items. In addition, some items must be specified while keeping in mind the limitations of the HDI system.

The CIFS administrator must keep the following in mind when managing CIFS shares from MMC:

- The CIFS administrator cannot perform any operations on special Windows folders, such as `IPC$` and `ADMIN$`, displayed in the list of CIFS shares.
- If a CIFS share is deleted, share level ACL settings are also deleted.

Viewing a list of CIFS shares

When viewing a list of CIFS shares in the HDI system from MMC, some items are not available. The following table describes the displayed items and their availability in an HDI system.

Table 9-2 Items that can be viewed in the CIFS share list and their reliability in an HDI system

Item	Reliability	Explanation
Share Name	Y	The CIFS share name is displayed. Example: <code>share1</code>

Item	Reliability	Explanation
Folder Path	Y	The CIFS share path is displayed. Example: C:\mnt\fs01\share1
Type	N	One of the following network connection types is displayed, although the value is not correct: <ul style="list-style-type: none"> • Windows • Macintosh • NetWare
# Client Connections	Y	The number of clients connected to the CIFS share is displayed.
Description	Y	The description of the CIFS share specified when the CIFS share was created is displayed. Example: share1

Legend

Y: Available, N: Not available

Creating a CIFS share

When you use MMC to create a CIFS share, the HDI system restrictions apply to the characters you can use. The following describes the items to be specified in MMC.

Table 9-3 Items specified in MMC when creating a CIFS share

Item name	Description
Folder path	<p>Specify an absolute path beginning with <code>c:</code> for the CIFS share to be created.</p> <p>Use no more than 249 characters, not including the initial <code>c:</code>.</p> <p>After <code>c:</code>, you can use any alphanumeric character, exclamation mark (!), hash mark (#), dollar sign (\$), percent sign (%), ampersand (&), single quotation mark ('), left parenthesis ((), right parenthesis ()), plus sign (+), comma (,), hyphen (-), period (.), forward slash (/), semicolon (;), equal sign (=), at mark (@), left square bracket ([), right square bracket (]), caret (^), underscore (_), grave accent mark (`), left curly bracket ({), right curly bracket (}), tilde (~), or space. You can also specify multi-byte characters. Note that a forward slash or space specified at the end of the string will be removed.</p> <p>If you use a percent sign (%) in the absolute path, make sure the percent sign is not used in any of the following combinations:</p> <p><code>%D, %G, %H, %I, %L, %M, %N, %P, %R, %S, %T, %U, %V, %a, %d, %g, %h, %i, %m, %p, %u, %v, %w, %\$</code></p> <p>You cannot specify a path that contains symbolic links. Note that the directory names <code>.conflict</code>, <code>.conflict_longpath</code>, <code>.snaps</code>, <code>.history</code> and <code>.lost+found</code> cannot be specified, and the directory</p>

Item name	Description
	names <code>.arc</code> , <code>.system_gi</code> , <code>.system_reorganize</code> , and <code>lost+found</code> cannot be specified directly under a file system.
Share name	<p>Specify the name of the CIFS share to be created.</p> <p>Use no more than 80 characters.</p> <p>You can use any alphanumeric character, exclamation mark (!), hash mark (#), dollar sign (\$), percent sign (%), ampersand (&), single quotation mark ('), left parenthesis ((), right parenthesis ()), plus sign (+), comma (,), hyphen (-), period (.), semicolon (;), equal sign (=), at mark (@), left square bracket ([), right square bracket (]), caret (^), underscore (_), grave accent mark (`), left curly bracket ({), right curly bracket (}), tilde (~), or space. You can also specify multi-byte characters. However, the string cannot contain only a dollar sign or periods (e.g., \$, ., or ..) and cannot end with a period (e.g., <code>Abc.</code>). If the string ends with a dollar sign, you cannot specify a period just before that dollar sign (e.g., <code>Abc.\$</code>). The space specified at the end of the string will be removed.</p> <p>If you use a percent sign (%) in the share name, make sure the percent sign is not used in any of the following combinations:</p> <p><code>%D, %G, %H, %I, %L, %M, %N, %P, %R, %S, %T, %U, %V, %a, %d, %g, %h, %i, %m, %p, %u, %v, %w, %\$</code></p> <p>In addition, the CIFS share name cannot be <code>global</code>, <code>homes</code>, <code>printers</code>, <code>admin\$</code>, <code>c\$</code>, <code>global\$</code>, <code>homes\$</code>, <code>ipc\$</code>, or <code>printers\$</code>.</p> <p>This specification is not case sensitive. Specify a name that is unique to the node.</p>
Description	<p>Specify a description for the CIFS share to be created.</p> <p>Use no more than 256 characters.</p> <p>You can specify alphanumeric characters, exclamation marks (!), hash marks (#), dollar signs (\$), ampersands (&), single quotation marks ('), left parentheses ((), right parentheses ()), asterisks (*), plus signs (+), commas (,), hyphens (-), periods (.), forward slashes (/), colons (:), left angle brackets (<), right angle brackets (>), question marks (?), at marks (@), left square brackets ([), backslashes (\), right square brackets (]), carets (^), underscores (_), grave accent marks (`), left curly brackets ({), vertical bars (), right curly brackets (}), and tildes (~). You can also specify spaces except at the beginning or end of a character string. You cannot specify a backslash (\) at the end of a character string.</p> <p>You can also specify multi-byte characters.</p>

Changing CIFS share information

When you use MMC to change CIFS share information, HDI system restrictions apply to the values specified for some items. The following describes the items to be specified in MMC.

Table 9-4 Items specified for changing CIFS share information

Item name	Description
Description	<p>Specify a description for the CIFS share to be created.</p> <p>Use no more than 256 characters.</p> <p>You can specify alphanumeric characters, exclamation marks (!), hash marks (#), dollar signs (\$), ampersands (&), single quotation marks ('), left parentheses ((), right parentheses ()), asterisks (*), plus signs (+), commas (,), hyphens (-), periods (.), forward slashes (/), colons (:), left angle brackets (<), right angle brackets (>), question marks (?), at marks (@), left square brackets ([), backslashes (\), right square brackets (]), carets (^), underscores (_), grave accent marks (`), left curly brackets ({), vertical bars (), right curly brackets (}), and tildes (~). You can also specify spaces except at the beginning or end of a character string. You cannot specify a backslash (\) at the end of a character string.</p> <p>You can also specify multi-byte characters.</p>
Restriction on the number of users	<p>Specify the upper limit number of users who can connect to a CIFS share.</p> <p>Note that the number of connected users cannot exceed the maximum CIFS client connections in the HDI system. For details about the maximum number of CIFS client connections possible in the HDI system, see Table 7-2 Maximum number of CIFS client connections and the maximum number of CIFS shares on page 7-4.</p>
Offline Settings	<p>If there are files and programs available offline, one of the following options is selected to determine what items are made available to offline users:</p> <ul style="list-style-type: none"> • Only the files and programs that users specify are available offline • No files or programs from the shared folder are available offline • All files and programs that users open from the shared folder are automatically available offline

Session management from MMC

The CIFS administrator can use MMC to view a list of sessions for users who are accessing a CIFS share and to disconnect those sessions.

This section contains notes on managing, from MMC, sessions for users that are accessing CIFS shares in an HDI system.

Viewing a list of sessions

When viewing a list of sessions in which users are accessing a CIFS share in an HDI system, some items are not available. The following table describes the displayed items and their availability in an HDI system.

Table 9-5 Items that can be viewed in the list of sessions and their availability in an HDI system

Item	Availability	Explanation
User	Y	The names of users connected to a CIFS share are displayed. Example: group01\administrator
Computer	Y	The computer names of users connected to a CIFS share are displayed. Example: adam
Type	N	One of the following network connection types is displayed, although the value is not correct: <ul style="list-style-type: none"> • Windows • Macintosh • NetWare
# Open Files	Y	The number of files that the CIFS client currently has open in the CIFS shares is displayed.
Connected Time	N	The correct value is not displayed (0 is always displayed).
Idle Time	N	The correct value is not displayed (0 is always displayed).
Guest	N	The correct value is not displayed (No is always displayed).

Legend

Y: Available, N: Not available

Note that if Trend Micro ServerProtect is set to be used, sessions resulting from accesses from a scan server are also displayed.

Closing sessions

Note the following when you close sessions from MMC.

- When you specify a session to be disconnected, if there are multiple sessions that have the same user name and computer name as the specified session, all applicable sessions will be disconnected.
- When you close sessions, the data you are using might be lost. Before closing any sessions, contact all connected users.
- When a CIFS share is created or updated from MMC, the session is automatically disconnected. For this reason, if you use that session to access a file from the client machine on which you operated MMC, an error might occur. If an error occurs, disconnect the session, and then access the file again.
- If a session resulting from access from a scan server is closed while a virus scan is in progress, a scan error might occur.

Managing open files from MMC

The CIFS administrator can use MMC to view a list of files on a CIFS share that have been opened by a CIFS client and close the files.

This section provides notes on managing the CIFS share files, from MMC, that have been opened in an HDI system by CIFS clients.

List of open files

Some of the items in the list of files opened by CIFS clients on a CIFS share cannot be used in an HDI system. The following table lists the displayed items and their availability in an HDI system.

Table 9-6 Items displayed in the list of open files in an HDI system and their availability

Item	Available	Explanation
Open File	Y	Displays the names of the files opened by CIFS clients on the CIFS shares. Pipes that have a name attached are also included. A maximum of 260 characters are displayed for each file name. Example: C:\mnt\share\file.txt
Accessed By	Y	Displays the names of the users accessing the open files. Example: group01\user01
Type	N	One of the following network connection types is displayed. Note that the value is actually incorrect: <ul style="list-style-type: none">• Windows• Macintosh• NetWare
# Locks	Y	Displays the number of locks on each open file.
Open Mode	Y	Displays one of the following access permissions, which has been granted to the CIFS clients for each open file: <ul style="list-style-type: none">• Read• Write• Write+Read• no access

Legend:

Y: Available, N: Not available

Closing open files

You should note the following points when closing a file from MMC.

- The file will be closed forcibly without the CIFS client being notified, which means that any data that has not been saved might be lost. Before closing a file, you should contact the user who is working on the file.
- Pipes with names attached cannot be closed.
- When a directory is closed, the files in the directory are not closed.
- Closing a large volume of open files from MMC puts a heavy load on the HDI system.

Share-level ACLs

The CIFS administrator can use MMC to set up share level ACLs. Share level ACLs that are set up are applied to all files and subfolders in the shared folder.

Share level ACLs are set up for CIFS shares. They are not set up for individual directories and files in an HDI system like Advanced ACLs and Classic ACLs are. In contrast to the share level ACLs, Advanced ACLs and Classic ACLs are called file level ACLs.

When both a file level ACL and a share level ACL are set up, the share level ACL is evaluated before the file level ACL. For example, if read-only permission is specified in the share level ACL and full control of that file is specified in the file level ACL, the actual permission will be read-only.

The following table lists the share level ACL information that the CIFS administrator can specify.

Table 9-7 Share level ACL

Item	Description
Access permissions that can be specified	<ul style="list-style-type: none"> • Full Control • Change • Read
ACE type	Allow or Deny
Maximum number of ACEs	1,820 entries per share

When setting up share level ACLs, the number of specifiable ACEs and CIFS shares has a limit for each resource group. Specify a value that meets the following condition.

$$65,536 > (\uparrow ((36 \times \text{number-of-ACEs} + 320) \times 10) \uparrow) \times \uparrow (\text{number-of-shares} / 10) \uparrow / 1024$$

Legend:

$\uparrow ((36 \times \text{number-of-ACEs} + 320) \times 10) \uparrow$: The value of the result of the calculation $((36 \times \text{number-of-ACEs} + 320) \times 10)$ rounded up to the nearest 8,192-byte unit.

$\uparrow (\text{number-of-shares} / 10) \uparrow$: The value of the result of the calculation $(\text{number-of-shares} / 10)$ rounded up to the nearest integer value.

The following shows the number of CIFS shares and a rough estimation of the specifiable number of ACEs for the shared directory.

Table 9-8 Rough estimation of the specifiable number of ACEs for CIFS shares:

Number of CIFS shares	Number of ACEs for the shared directory
1,000	1,820
7,500	210

You can specify *Full Control*, *Change*, and *Read* as access permissions in a share level ACL. The following table lists the access permissions and the available operations for a CIFS share.

Table 9-9 Access permissions specified in a share level ACL and the available operations for a CIFS share

Operation for a CIFS share	Access permissions in a share level ACL		
	Full Control	Change	Read
Displaying file names and subfolder names	Yes	Yes	Yes
Moving files to a subfolder	Yes	Yes	Yes
Displaying file contents and executing programs	Yes	Yes	Yes
Adding files and subfolders to a shared folder	Yes	Yes	No
Changing file data	Yes	Yes	No
Deleting subfolders and files	Yes	Yes	No
Changing access permissions	Yes	No	No
Acquiring ownership	Yes	No	No

Legend:

Yes: Possible, No: Not possible

Note the following when you set a share-level ACL:

- You cannot set a share-level ACL from File Services Manager.
- A share-level ACL is effective only for CIFS shares. The access restrictions of a share-level ACL are not applied to accesses from the NFS service.
- If more than one type of access permission can apply to a user, the logical OR of each access permission is used. For example, when user *A*, who has the Read permission, belongs to group *B*, which has the Change permission, the Change permission is applied to user *A*.
- If you specify a **Deny** entry and an **Allow** entry at the same time, the **Deny** entry has priority.
- When you change a share-level ACL, the changed ACL is not applied to users who are already connected to the target CIFS share. Before

changing a share-level ACL, the CIFS administrator must make sure that no users are connected to the target CIFS share.

- If access control (`read only`, `read list`, or `write list`) is specified for a CIFS share in addition to a share level ACL, the stricter of the two is applied. For example, if full control is specified in the share level ACL and `read only` is specified in the access control, the target CIFS share can only be read.

The following table lists the access permissions that are applied when a share level ACL and access control are set up for a CIFS share.

Table 9-10 Access permissions applied when a share level ACL and access control are set up for a CIFS share

Share level ACL	Access control			
	read only		read list	write list
	yes	no		
Full Control	RO	RW	RO	RW
Change	RO	RW	RO	RW
Read	RO	RO	RO	RO

Legend:

RO: Read-only, RW: Read and write

- If you specify only `Read` for a share-level ACL, the target CIFS shares are accessed as read-only shares. Similarly, the CIFS administrator can only access the shares as read-only.

Notes on using MMC

This subsection provides notes on using MMC. The screenshots appearing in this subsection are based on MMC 3.0 in Windows Server 2012 R2.

- About the **Browse For Folder** window

When you add a share, you can choose a path to the folder to be shared in the **Browse For Folder** window (see [Figure 9-1 Browse For Folder window on page 9-13](#)).

In the **Browse For Folder** window, a file system is displayed as a folder under the C\$ folder. Only mounted file systems are displayed.

If a node is failed over, the other (active) node takes both resource groups under its control and the file systems on the failed-over node become viewable on the active node. However, in the event that this happens, please be aware that file shares cannot be created for the file system on the node that failed over, but operations like selecting the file system and creating folders within the file system on that node can still be done.

If you select a file system and a folder on the node that failed over and attempt to create a file share, an error will occur (For details about the

various errors that could occur, see [Errors occurring when a share is added on page A-9](#) in this manual).

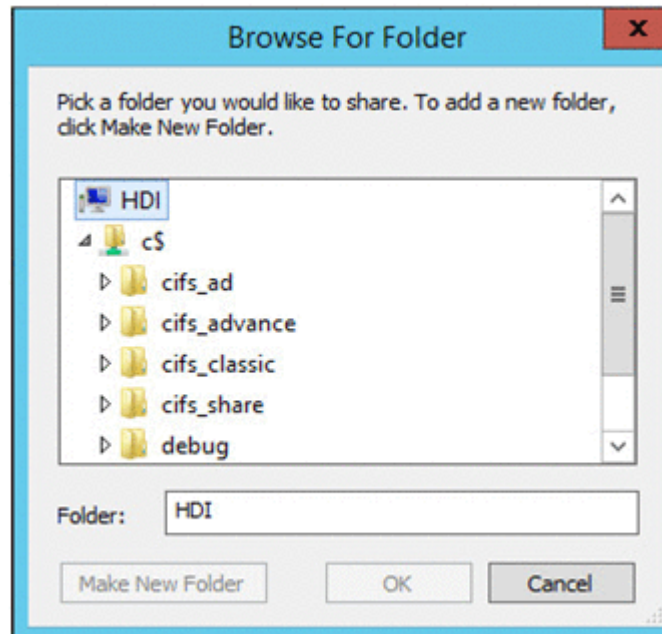


Figure 9-1 Browse For Folder window

- Disconnecting a user session
By selecting a user, and then clicking **Close Session** on the **Action** menu, you can close only the session of that user (see [Figure 9-2 Console window with the Action menu expanded - 1 on page 9-13](#)). Note that if there are multiple sessions that have the same user name and computer name as the target user, those sessions will all be disconnected at the same time.




Figure 9-2 Console window with the Action menu expanded - 1

- Disconnecting all sessions
By selecting **Sessions** in the tree, and then choosing **Disconnect All Sessions** from the **Action** menu, you can close all connected sessions at the same time (see [Figure 9-3 Console window with the Action menu expanded - 2 on page 9-14](#)). If multiple users that have the same user name and computer name are connected to the sessions, the sessions are closed successfully, but the error window shown in [Figure A-4 Window](#)

appearing when a session is not closed on page A-13 appears repeatedly.

If the same error window appears but the sessions are not closed, see [Disconnecting a session fails due to access denial on page A-13](#).

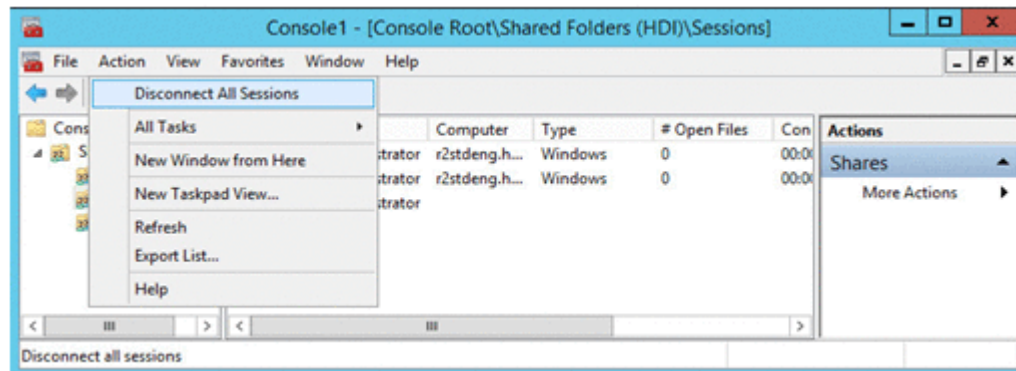


Figure 9-3 Console window with the Action menu expanded - 2

- Closing a file
By selecting a file and then choosing **Close Open File** from the **Action** menu, you can close the selected file (see [Figure 9-4 Sample window for closing a selected file on page 9-14](#)). Note that the selected file will be closed without the CIFS client being notified.

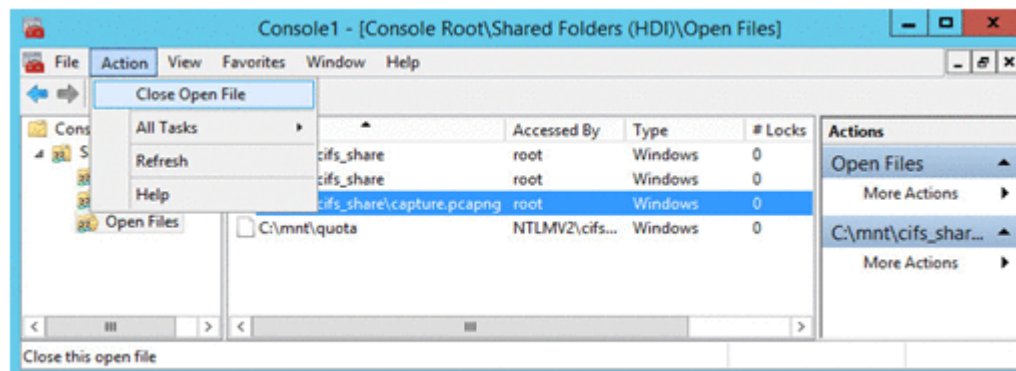


Figure 9-4 Sample window for closing a selected file

- Closing all files
By selecting **Open Files** in the tree, and then choosing **Disconnect All Open Files** from the **Action** menu, you can close all of the files in the CIFS shares that have been opened by CIFS clients (see [Figure 9-5 Sample windows for closing all of the files on page 9-15](#)). Note that the files will be closed without the CIFS clients being notified.

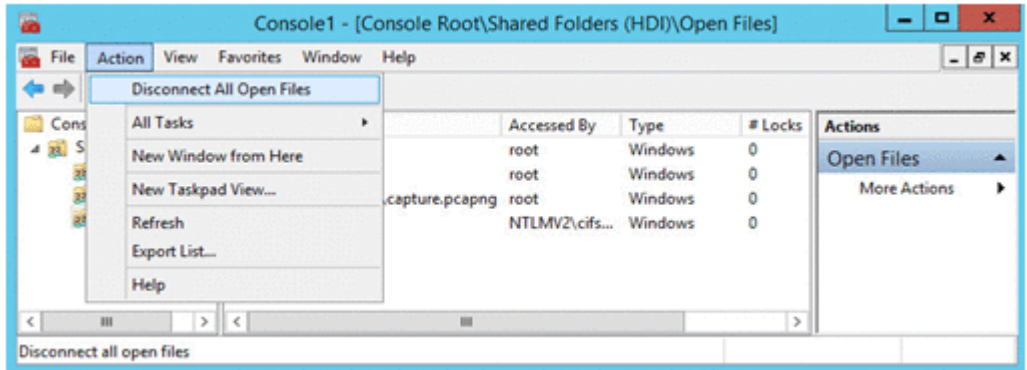


Figure 9-5 Sample windows for closing all of the files

- MMC version differences
When creating CIFS shares from MMC, the default access permissions differ depending on MMC version as shown in the following table.

Table 9-11 Default access permissions for different MMC versions

MMC version	Default access permissions
1.2	Full control permission is granted to all users.
2.0 or 3.0	Read-only permission is granted to all users.

For MMC 2.0 or 3.0, if you want to grant full control or modify permission to users for a CIFS share, use the wizard to open the access permissions setting window, select the **Use custom share and folder permissions** radio button, and then click the **Customize** button to grant **Full Control** or **Modify** permission to the users.

When you use MMC 3.0 to delete CIFS shares, the message shown below will appear. The CIFS shares will be deleted after you click the **Yes** button.

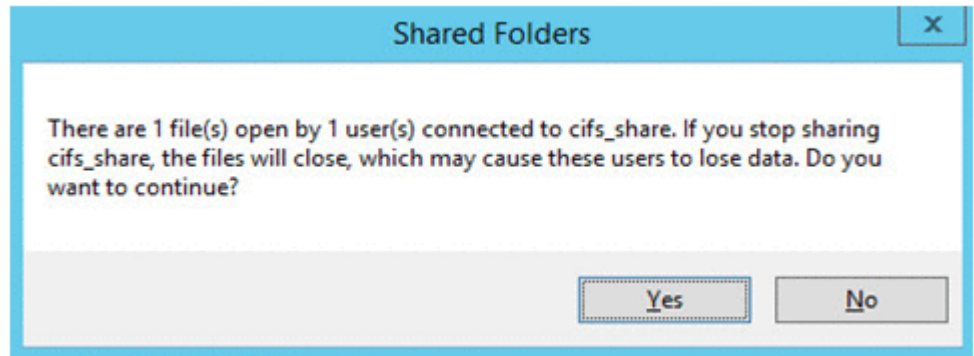


Figure 9-6 Message that appears when MMC 3.0 is used to delete CIFS shares

Note that when you delete CIFS shares, data that is being modified by other users might be lost. Before deleting a CIFS share, be sure to inform the users who can access the share.

Making Past Versions of Files Available by Using Volume Shadow Copy Service

This chapter explains how to make past versions of files that have been migrated to an HCP system available to CIFS clients by using Volume Shadow Copy Service.

For details about making past versions of files migrated to an HCP system available, see the *Installation and Configuration Guide*.

- [Overview of Volume Shadow Copy Service](#)
- [Compatible CIFS client platforms for Volume Shadow Copy Service](#)
- [Notes on using Volume Shadow Copy Service on a CIFS client](#)

Overview of Volume Shadow Copy Service

Volume Shadow Copy Service can be used to make past versions of files migrated to an HCP system available to CIFS clients.

When a CIFS client opens the Properties dialog box of a file or folder in the file system, past versions of that file or folder are listed in the **Previous Versions** tab. By using this tab, the CIFS client can view the contents of a past version, copy a past version to another folder, or restore a file or folder from a past version.

The following figure shows an example of the **Previous Versions** tab.

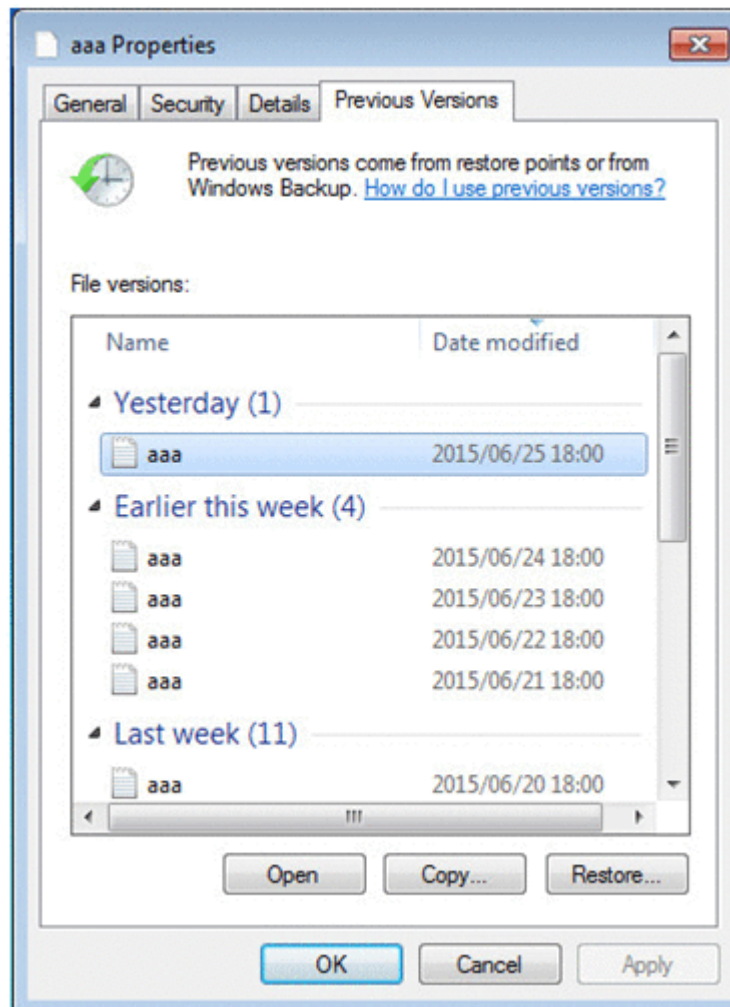


Figure 10-1 Previous Versions tab in the Properties dialog box for a file or folder

The **Previous Versions** tab can display past versions of a file or folder from the last seven days. Versions older than those that are displayed in this tab are viewed in the `.history` directory.

Compatible CIFS client platforms for Volume Shadow Copy Service

Among the CIFS client platforms supported by an HDI system, the following platforms support Volume Shadow Copy Service:

- Windows 7
- Windows 8
- Windows 10
- Windows Server 2008
- Windows Server 2012
- Windows Server 2016

Notes on using Volume Shadow Copy Service on a CIFS client

When you use Volume Shadow Copy Service to view past versions of files migrated to an HCP system from a CIFS client, the following notes apply:

- To access past versions of a file or folder via Volume Shadow Copy Service, use the **Previous Versions** tab of the Properties dialog box for the file or folder.
- The past versions of a file listed in the **Previous Versions** tab might differ depending on the platform of the CIFS client.
- If files have the offline attribute, view them in the `.history` directory because the **Previous Versions** tab does not display them.
- For files and directories that existed, were removed or deleted, and then created again, the past versions of the files or directories might not be displayed in the **Previous Versions** tab. In this case, use the `.history` directory to view the past versions.
- If you are accessing the past versions by using the Volume Shadow Copy Service, additional information required to access the past versions is added to the path name. Therefore, the maximum lengths of accessible file paths and directory paths will be 25 characters shorter than the maximum lengths when you access the past versions from a client. If the path length containing the additional information exceeds the maximum length of paths that can be accessed from a client, you will no longer be able to access the past versions from a client. The maximum length of paths that can be accessed from a client varies depending on the client you are using and the environment.
- When you cannot access to the past versions by using Volume Shadow Copy Service, due to the restriction of file name and/or folder path length, please make sure to shorten the path length by mapping the folder, locates just above where the target file or folder you are looking at, to the network drive, then access to the target file or the folder.
- The operation performed by a CIFS client to display the **Previous Versions** tab places load on the system. If this operation is performed

either frequently or by multiple clients at the same time, file-system access performance might be affected by up to around 10%. We therefore recommend you perform this operation at a time when system load is low. And, in case the system is overwhelmed by User I/O almost hitting to its system limitation, the user may experience roughly 35% access performance degradation while the operation above.

- After performing a copy or restoration from the **Previous Versions** tab of the Properties dialog box, to perform the copy or restoration again, close and then re-open the dialog box. You cannot perform the copy or restoration again without closing the Properties dialog box.
- When a file or folder is copied or restored by using the **Previous Versions** tab from a CIFS client, the following attributes are not inherited:

- Created date and time
- Accessed date and time
- ACL
- Owner

- When specifying a file or folder name in a file share, do not use the following format:

@GMT-*nnnn.nn.nn-*nn.nn.nn** (where *n* is a number)

If you assign a name in the above format to a file or folder, you might not be able to access that file or folder, or its past versions.

- The operations that can be performed from a CIFS client depend on the file type, as shown in the table below.

Table 10-1 Restrictions on operations from the CIFS client

Operation	Non-WORM file	WORM file#	Stub file
View a list of the target files in the Previous Versions tab	Y	Y	N
Open button	Y	Y	N
Copy button	Y	Y	N
Restore button	Y	N	N

Legend:

Y: Can be executed

N: Cannot be executed

#

Includes WORM files whose retention period has passed.

- If an error occurs in communication with the HCP system while a file is being restored by using the **Previous Versions** tab, restoration fails. As a result, the file might be deleted from the file system. If the file is deleted, correct the communication error, and then use either of the following methods to restore the file:

- Copy the file from a past-version directory in the `.history` directory.
- Open the Properties dialog box of the parent folder, and then, in the **Previous Versions** tab, select and open the past version of the folder that you want to use for restoration. Then, copy the file that you opened.

CIFS Client Platforms

This chapter contains notes on the differences of using various CIFS client platforms.

- [Notes common to all supported types of Windows](#)
- [Notes for Windows Server 2008](#)
- [Notes for Windows 7](#)
- [Notes for Windows 8](#)
- [Notes for Windows 10 or Windows Server 2016](#)
- [Notes for Windows Server 2012](#)
- [Notes for Mac OS X](#)

Notes common to all supported types of Windows

When a CIFS client accesses a CIFS share for the first time after logging on to Windows, the user name and the password used for logging on to Windows are used to send an authentication request to an HDI node. This means that when you have permitted access by a guest account, the user who has logged on to Windows might access the CIFS share as a guest account without being requested to input a user name and a password. This might occur when the user is specified in **Mapping to guest account** in the **CIFS Service Management** page (**Setting Type:** *Security*) of the **Access Protocol Configuration** dialog box. Therefore, be careful when you permit access to a guest account.

In the CIFS service authentication mode of Active Directory authentication, if an authentication attempt fails for the user name and the password used to log on to Windows, you will be asked to input a user name and a password. When access by a guest account is allowed and entry of a user name and password results in an authentication failure, access will automatically be granted using the guest account. Please keep this in mind.

Notes for Windows Server 2008

This section contains notes on when the CIFS client is Windows Server 2008.

Files and folders in shared directories

This subsection contains notes on files and folders to be created in a shared directory from the client.

When adding an ACL

If the following conditions are met during an attempt to add an ACL to a file or directory in the HDI share, you cannot add the ACL because the system cannot refer to the local user or group:

- You can use one of the following methods to log on to the client:
 - With the user account control (UAC) enabled, you can log on as either of the following local accounts:
 - Administrator (a built-in account)
 - Non-administrative user account (a local account on the client machine)
 - With the user account control (UAC) disabled, you can log on to the client using a local account.
- The authentication method is one of the following:
 - Local authentication
 - NT domain authentication without user mappings
 - Active Directory authentication without user mappings

When using quotas

When a client moves or pastes a file into a share by using Explorer, if the amount of used blocks or the number of inodes has exceeded the software limit specified on the HDI system, the client operation fails even if the hardware limit is not exceeded. Note that you can use the `COPY` or `XCOPY` command to move or paste files into a share, if the hardware limit is not exceeded.

When enabling offline files

When you create Microsoft Office files in a share in which you have enabled offline files from the client, you might be unable to correctly save the files. Therefore, disable offline files in shares when you use Microsoft Office on Windows Server 2008.

When using a network drive

When you use a network drive to display the properties of a file or directory whose path name is close to the maximum length, the security tab will not be displayed, which prevents you from viewing or configuring the ACL. If the security tab is not displayed in the properties, temporarily make the file or directory name shorter.

The security tab is not displayed if all of the following conditions are met:

- You use a network drive.
- The path name of the target file or directory on the network drive exceeds 259 characters when represented in the UNC format (`\\server-name\share-name\...`).

When using MMC

This subsection contains notes for when you use MMC from Windows Server 2008 to manage a CIFS share.

Logging on to Windows

Log on to Windows using one of the methods shown below. HDI will deny access by any other method.

- Use a domain account for logging on.
- Use an `Administrator` account for logging on.
- With the user account control (UAC) disabled on the client, use an administrator account except for `Administrator` for logging on.

Share-level ACLs

If the following conditions are met during an attempt to add an entry to a share-level ACL, you cannot add an entry to the share-level ACL because the system cannot refer to the local user or group:

- You can use one of the following methods to log on to the client:

- With the user account control (UAC) enabled, you can log on as either of the following local accounts:
 - Administrator (a built-in account)
 - Non-administrative user account (a local account on the client's computer)
- With the user account control (UAC) disabled, you log on to the client using a local account.
- The authentication method is one of the following:
 - Local authentication
 - NT domain authentication without user mappings
 - Active Directory authentication without user mappings

When using a large 1-MB MTU

When the system is first installed, a large 1-MB MTU is disabled. For details about how to enable a large 1-MB MTU on the CIFS client, contact Microsoft Support.

Notes when attempting to access the CIFS service

To access a CIFS share by using SMB2, you must first perform the procedure described in article KB978625 of the Microsoft Knowledge Base. Read the article, and then contact Microsoft Support. If you attempt to access a CIFS share by using SMB2 from a Windows client without first performing the procedure, a STOP error might occur and one of the following error messages might be displayed:

```
- STOP: 0x00000027 (parameter1, parameter2, parameter3,
parameter4)
```

```
- mrxsmb20.sys - Address parameter1 base at parameter2, Datestamp
parameter3
```

Notes when you are accessing the CIFS service

If a failover or failback occurs while a client is accessing the CIFS service, an error might occur the next time the CIFS service is accessed. If this happens, try to access the CIFS service one more time.

Notes for Windows 7

This section describes notes on when the CIFS client is Windows 7.

Files and folders in shared directories

This subsection describes notes on files and folders to be created in a shared directory from the client.

When adding an ACL

If the following conditions are met during an attempt to add an ACL to a file or directory in the HDI share, you cannot add the ACL because the system cannot refer to the local user or group:

- You can use one of the following methods to log on to the client:
 - With the user account control (UAC) enabled, you can log on as either of the following local accounts:
 - Administrator (a built-in account)
 - Non-administrative user account (a local account on the client's computer)
 - With the user account control (UAC) disabled, you can log on to the client using a local account.
- The authentication method is one of the following:
 - Local authentication
 - NT domain authentication and no user mapping used
 - Active Directory authentication and no user mapping used

When using quotas

When a client moves or pastes a file into a share by using Explorer, if the amount of used blocks or the number of inodes has exceeded the software limit specified on the HDI system, the client operation fails even if the hardware limit is not exceeded. Note that you can use the `COPY` or `XCOPY` command to move or paste files into a share, if the hardware limit is not exceeded.

When using network drives

When you use a network drive to display the properties of a file or directory whose path name is close to the maximum length, the security tab will not be displayed, which prevents you from viewing or configuring the ACL. If the Security page is not displayed in the Properties window, temporarily shorten the file directory name.

The security tab is not displayed if all of the following conditions are met:

- You use a network drive.
- The path name of the target file or directory on the network drive exceeds 259 characters when represented in the UNC format (`\\server-name\share-name\...`).

When enabling offline files

When you create Microsoft Office files in a share in which you have enabled offline files from the client, you might be unable to correctly save the files. Therefore, disable offline files in shares when you use Microsoft Office on Windows 7.

When using MMC

This subsection contains notes on when you use MMC from Windows 7 to manage a CIFS share.

Logging on to Windows

Log on to Windows using one of the methods shown below. If you use another method, access to HDI will be denied.

- Use a domain account for logging on.
- Use an `Administrator` account for logging on.
- With the user account control (UAC) disabled on the client, use an administrator account except for `Administrator` to log on.

Share-level ACLs

If the following conditions are met during an attempt to add an entry to a share-level ACL, you cannot add an entry to the share-level ACL because the system cannot refer to the local user or group:

- You can use one of the following methods to log on to the client:
 - With the user account control (UAC) enabled, you can log on as either of the following local accounts:
 - `Administrator` (a built-in account)
 - Non-administrative user account (a local account on the client's computer)
 - With the user account control (UAC) disabled, you can log on to the client using a local account.
- The authentication method is one of the following:
 - Local authentication
 - NT domain authentication without user mappings
 - Active Directory authentication without user mappings

When using a large 1-MB MTU

When the system is first installed, a large 1-MB MTU is disabled. For details about how to enable a large 1-MB MTU on the CIFS client, contact Microsoft Support.

Notes for Windows 8

This section describes notes on when the CIFS client is Windows 8.

Files and folders in shared directories

This subsection describes notes on files and folders to be created in a shared directory from the client.

When adding an ACL

If the following conditions are met during an attempt to add an ACL to a file or directory in the HDI share, you cannot add the ACL because the system cannot refer to the local user or group:

- You can use one of the following methods to log on to the client:
 - With the user account control (UAC) enabled, you can log on as either of the following local accounts:
 - Administrator (a built-in account)
 - Non-administrative user account (a local account on the client's computer)
 - With the user account control (UAC) disabled, you can log on to the client using a local account.
- The authentication method is one of the following:
 - Local authentication
 - NT domain authentication and no user mapping used
 - Active Directory authentication and no user mapping used

When using quotas

When a client moves or pastes a file into a share by using Explorer, if the amount of used blocks or the number of inodes has exceeded the software limit specified on the HDI system, the client operation fails even if the hardware limit is not exceeded. Note that you can use the `COPY` or `XCOPY` command to move or paste files into a share, if the hardware limit is not exceeded.

When enabling offline files

When you create Microsoft Office files in a share in which you have enabled offline files from the client, you might be unable to correctly save the files. Therefore, disable offline files in shares when you use Microsoft Office on Windows 8.

When using MMC

This subsection contains notes on when you use MMC from Windows 8 to manage a CIFS share.

Logging on to Windows

Log on to Windows using one of the methods shown below. If you use another method, access to HDI will be denied.

- Use a domain account for logging on.
- Use an `Administrator` account for logging on.
- With the user account control (UAC) disabled on the client, use an administrator account except for `Administrator` to log on.

Share-level ACLs

If the following conditions are met during an attempt to add an entry to a share-level ACL, you cannot add an entry to the share-level ACL because the system cannot refer to the local user or group:

- You can use one of the following methods to log on to the client:
 - With the user account control (UAC) enabled, you can log on as either of the following local accounts:
 - `Administrator` (a built-in account)
 - Non-administrative user account (a local account on the client's computer)
 - With the user account control (UAC) disabled, you can log on to the client using a local account.
- The authentication method is one of the following:
 - Local authentication
 - NT domain authentication without user mappings
 - Active Directory authentication without user mappings

Notes for Windows 10 or Windows Server 2016

This section describes notes on when the CIFS client is Windows 10.

Files and folders in shared directories

This subsection describes notes on files and folders to be created in a shared directory from the client.

When adding an ACL

If the following conditions are met during an attempt to add an ACL to a file or directory in the HDI share, you cannot add the ACL because the system cannot refer to the local user or group:

- You can use one of the following methods to log on to the client:
 - With the user account control (UAC) enabled, you can log on as either of the following local accounts:
 - `Administrator` (a built-in account)
 - Non-administrative user account (a local account on the client's computer)

- With the user account control (UAC) disabled, you can log on to the client using a local account.
- The authentication method is one of the following:
 - Local authentication
 - NT domain authentication and no user mapping used
 - Active Directory authentication and no user mapping used

When using quotas

When a client moves or pastes a file into a share by using Explorer, if the amount of used blocks or the number of inodes has exceeded the software limit specified on the HDI system, the client operation fails even if the hardware limit is not exceeded. Note that you can use the `COPY` or `XCOPY` command to move or paste files into a share, if the hardware limit is not exceeded.

When enabling offline files

When you create Microsoft Office files in a share in which you have enabled offline files from the client, you might be unable to correctly save the files. Therefore, disable offline files in shares when you use Microsoft Office on Windows 10.

When using MMC

This subsection contains notes on when you use MMC from Windows 10 to manage a CIFS share.

Logging on to Windows

Log on to Windows using one of the methods shown below. If you use another method, access to HDI will be denied.

- Use a domain account for logging on.
- Use an `Administrator` account for logging on.
- With the user account control (UAC) disabled on the client, use an administrator account except for `Administrator` to log on.

Share-level ACLs

If the following conditions are met during an attempt to add an entry to a share-level ACL, you cannot add an entry to the share-level ACL because the system cannot refer to the local user or group:

- You can use one of the following methods to log on to the client:
 - With the user account control (UAC) enabled, you can log on as either of the following local accounts:
 - `Administrator` (a built-in account)

- Non-administrative user account (a local account on the client's computer)
- o With the user account control (UAC) disabled, you can log on to the client using a local account.
- The authentication method is one of the following:
 - o Local authentication
 - o NT domain authentication without user mappings
 - o Active Directory authentication without user mappings

Notes on accessing the CIFS service

If a failover or failback occurs while a client is accessing the CIFS service, an error might occur the next time the CIFS service is accessed. If this happens, try to access the CIFS service one more time.

Connecting by using the SMB 1.0 protocol

If you are using Windows Server 2016, SMB signing is required by default to establish a CIFS connection. If the HDI system is set up to use SMB 1.0, SMB signing is not used by default and attempts to access CIFS resources will end in an error. In this case, change either of the following settings:

CIFS client setting

On the CIFS client, change the setting related to CIFS access using SMB signing.

On the client, open **Local Security Policy**. Then, from **Security Settings**, select **Local Policies** and then **Security Options**. Check the setting **Microsoft network client: Digitally sign communications (always)**. If the setting is enabled, disable it.

HDI system setting (if the CIFS client setting cannot be changed)

Run the `cifsoptset` command so that SMB signing is always used during SMB 1.0 communication to access CIFS resources in the HDI system. For details about how to change the setting related to SMB signing, see the *CLI Administrator's Guide*.

Notes for Windows Server 2012

This section contains notes on when the CIFS client is Windows Server 2012.

Files and folders in shared directories

This subsection contains notes on files and folders to be created in a shared directory from the client.

When adding an ACL

If the following conditions are met during an attempt to add an ACL to a file or directory in the HDI share, you cannot add the ACL because the system cannot refer to the local user or group:

- You can use one of the following methods to log on to the client:
 - With the user account control (UAC) enabled, you can log on as either of the following local accounts:
 - `Administrator` (a built-in account)
 - Non-administrative user account (a local account on the client machine)
 - With the user account control (UAC) disabled, you can log on to the client using a local account.
- The authentication method is one of the following:
 - Local authentication
 - NT domain authentication without user mappings
 - Active Directory authentication without user mappings

When using quotas

When a client moves or pastes a file into a share by using Explorer, if the amount of used blocks or the number of inodes has exceeded the software limit specified on the HDI system, the client operation fails even if the hardware limit is not exceeded. Note that you can use the `COPY` or `XCOPY` command to move or paste files into a share, if the hardware limit is not exceeded.

When using MMC

This subsection contains notes for when you use MMC from Windows Server 2012 to manage a CIFS share.

Logging on to Windows

Log on to Windows using one of the methods shown below. HDI will deny access by any other method.

- Use a domain account for logging on.
- Use an `Administrator` account for logging on.
- With the user account control (UAC) disabled on the client, use an administrator account except for `Administrator` for logging on.

Share-level ACLs

If the following conditions are met during an attempt to add an entry to a share-level ACL, you cannot add an entry to the share-level ACL because the system cannot refer to the local user or group:

- You can use one of the following methods to log on to the client:
 - With the user account control (UAC) enabled, you can log on as either of the following local accounts:
 - Administrator (a built-in account)
 - Non-administrative user account (a local account on the client's computer)
 - With the user account control (UAC) disabled, you log on to the client using a local account.
- The authentication method is one of the following:
 - Local authentication
 - NT domain authentication without user mappings
 - Active Directory authentication without user mappings

Notes on accessing the CIFS service

If a failover or failback occurs while a client is accessing the CIFS service, an error might occur the next time the CIFS service is accessed. If this happens, try to access the CIFS service one more time.

Notes for Mac OS X

This section contains notes for Mac OS X CIFS clients.

Support range

Note the following limitations on Mac OS X CIFS clients:

- MMC linkage cannot be performed.
- Past versions of files that used Volume Shadow Copy Service cannot be made available to everyone.
- DFSs (Distributed File Systems) cannot be used.
- CIFS-share update data cannot be cached locally.

Notes on file and directory names

The types of characters and the length of paths that can be used for file and directory names are limited to those that are supported by Windows. Therefore, the following characters and path names cannot be used when specifying file and directory names:

- Spaces at the end of file names
- Double quotation marks ("), asterisks (*), forward slashes (/), colons (:), left angle brackets (<), right angle brackets (>), question marks (?), backslashes (\), and vertical bars (|)
- File path names consisting of 256 or more characters

- Directory path names consisting of 260 or more characters
- Path names that are longer than what the application permits
For example: in Excel 2004, a file path name consisting of 219 or more single-byte characters

Notes on operations

The following are notes on operations for Mac OS X CIFS clients:

In addition, see the notes in [Notes for Mac OS X v10.9 on page 11-14](#) for Mac OS X v10.9 and the notes in [Notes for Mac OS X v10.10, v10.11, or macOS v10.12 on page 11-15](#) for Mac OS X v10.10, v10.11 or macOS v10.12.

- When a client moves or pastes a file into a share, if the amount of used blocks or the number of inodes has exceeded the software limit specified on the HDI system, the client operation fails even if the hardware limit is not exceeded.
- If a CIFS share is accessed or a file operation is performed on a CIFS share, hidden files might be created. These hidden files also use the user area on the file system and increase the amount of blocks and inodes that are used. As such, hidden files must be taken into account when estimating the size of a file system or when setting quotas for an HDI system.
- If a user account used to access a file share contains multi-byte characters, character codes are used to display the highest-level share name in **Finder**.
- Depending on the applications used on a client, if you update a file that is in a CIFS share, the ACL for the file might be changed back to its initial configuration. To avoid this, configure the ACLs of higher-level directories instead of configuring the ACLs of individual files.
- Depending on the applications used on a client, if you update a file that is in a CIFS share, the file owner might be changed to another user who has performed an operation on the file. Configure the ACLs of higher-level directories so that the access control is not dependent on file owners. For Classic ACL file systems, set ACLs for the file operator and the groups the file operator belongs to.
- If more than 128 ACEs are set for a file or folder, only the first 128 ACEs are displayed. Also, attempting to set more than 128 ACEs will result in an error. If this happens, access permissions for the file or folder are determined by the set ACEs only.
- To modify an ACL, add an HDI node and Mac OS X CIFS client to the Active Directory domain, and then use a domain user account to log in to the CIFS client.
- When an operation (for example, an ACL operation) is performed on a CIFS share from a client running Mac OS X (which is based on UNIX), depending on whether the UNIX client extended functionality for the CIFS service is being used, how the request from the client to the CIFS share is handled differs as described below.

When the UNIX client extended functionality of the CIFS service is being used:

Because the UNIX client extended functionality is being used to perform processing, we recommend using Classic ACL file systems, for which ACLs based on POSIX ACLs can be set, to handle requests from CIFS clients. In the case of Advanced ACL file systems (for which ACLs based on NTFS ACLs can be set), requests that cannot be processed by the UNIX client extended functionality are instead processed on the client side (like Windows requests).

When the UNIX client extended functionality of the CIFS service is not being used:

For both Classic ACL file systems and Advanced ACL file systems, requests are processed on the client side (like Windows requests).

Use the `cifsoptset` command to set whether to use the UNIX client extended functionality. Use the `cifsoptlist` command to check whether the UNIX client extended functionality is being used.

- Versions (a new function for documents in Mac OS X v10.7 or later) cannot be used.

When you perform an operation on a document, the following message might be displayed: The document "*file name*" is on a volume that does not support permanent version storage. This message can be ignored.

- If you want to add user and group ACLs when user mapping is not used for local authentication, use the `dirsetacl` command to specify ACL settings. If you want to add an ACL to a file, use the `dirsetacl` command to specify the ACE inheritance range for the relevant directory so that the required ACEs are inherited by the target file.
- In Mac OS X, even if you have write permission for a file, writing to the file might fail due to an application running on Mac OS X.

Therefore, when you perform an operation that requires updating any files on Mac OS X, we recommend that you configure the settings as follows:

- a. Grant the Full Control permission to the user performing the operation or to the group that the user belongs to, for the `.TemporaryItems` folder immediately under the CIFS share and all files and folders in the `.TemporaryItems` folder.
- b. Set the Delete permission for the target file of the operation, or set the Delete Subfolders and Files permission for the parent folder.
- c. Set access permissions for both the user performing the operation and the group to which the user belongs, and then make sure that the access permissions are inherited by higher-level folders.

Notes for Mac OS X v10.9

Mac OS v10.9 clients cannot access CIFS by using SMB 2.0, SMB 2.1, or SMB 3.0. You need to change the setting in HDI system or on Mac OS X v10.9 clients to access CIFS by using SMB 1.0.

If you change the setting in HDI system:

In the **CIFS Service Management** page (**Setting Type:** *Basic*), specify the setting so that SMB 1.0 is used for access from CIFS clients. For details on how to do this, see the *Administrator's Guide*.

If you change the setting on Mac OS X v10.9 clients:

Change the setting by using one of the following methods:

- From **Connect to Server** in **Finder**, connect the server by using `cifs://` instead of `smb://`.

- Add the following text to the `~/Library/Preferences/nsmb.conf` file. Then, from **Connect to Server** in **Finder**, connect to the server by using `smb://`.

```
[default]
smb_neg=smb1_only
```

Notes for Mac OS X v10.10, v10.11, or macOS v10.12

Note the following when you use Mac OS X v10.10, v10.11, or macOS v10.12:

- Mac OS X v10.10, v10.11, and macOS v10.12 support only SMB 2.0. To allow Mac OS X v10.10, v10.11, and macOS v10.12 clients to access CIFS, you need to change the setting in the HDI system.
In the **CIFS Service Management** page (**Setting Type:** *Basic*), specify the setting so that SMB 2.0 is used for access from CIFS clients. For details on how to do this, see the *Administrator's Guide*.
Note that on Mac OS X v10.10, v10.11, and macOS v10.12 clients, a minor version such as *x* in SMB2.*x* cannot be specified.
- Mac OS X v10.9 and earlier versions support only CIFS access that uses SMB 1.0. Therefore, to use Mac OS X v10.9 and earlier versions together with v10.10, v10.11, or macOS v10.12, you need to change the setting on v10.9 and earlier versions of CIFS clients to restrict the SMB version to 1.0 on each client. For details about how to change the setting, see [Notes for Mac OS X v10.9 on page 11-14](#).
If you upgrade Mac OS X v10.9 or an earlier version to v10.10, v10.11, or macOS v10.12, use the **CIFS Service Management** page (**Setting Type:** *Basic*) to configure the settings so that SMB 2.0 is used for access from CIFS clients. Then, cancel the setting that restricts the SMB version to 1.0 on each client.
- If any multi-byte characters are used for CIFS share name with Mac OS X v10.11, because of a matter of Mac client, connection from the Mac client to CIFS may be disabled. Avoid the use of multi-byte characters for share names.

Overview of the NFS Service

NFS clients can access data via the NFS service in an HDI system. This chapter provides an overview of using the NFS service.

- [Overview of using the NFS service](#)

Overview of using the NFS service

If the system administrator creates NFS shares for file systems and directories, NFS clients can access data in the storage systems via a network. The following figure shows how NFS clients can access data in the file system.




Figure 12-1 Flow of NFS clients accessing data

File systems provided by HDI systems support the NFSv2, NFSv3, and NFSv4 protocols. You can use the NFS service configuration definition to specify which versions of the NFS protocol will be supported. For details about viewing and changing the NFS service configuration, see the *Administrator's Guide*.

NFS clients access file systems by using the NFS protocols that have been specified when the file system is mounted.

In an HDI system, NFS client user authentication can be performed by using the UNIX (AUTH_SYS) or Kerberos authentication method. For details on NFS client user authentication, see [Chapter 16, User Authentication for NFS Clients on page 16-1](#).

You must perform certain procedures before using the NFSv4 protocol, such as applying an appropriate patch.

As such, if you do not need to use the functions provided by the NFSv4 protocol, we recommend that you use either the NFSv2 or NFSv3 protocol.

System Configuration When the NFS Service Is Used

This chapter describes what kinds of HDI system configurations are required for using the NFS service and also the environments in which the NFS service can run.

- [Products supported by the NFS service](#)
- [Network configurations](#)
- [Configuring an NFS environment when Kerberos authentication and an NFSv4 domain configuration are used](#)

Products supported by the NFS service

This section describes the products supported by the NFS service.

NFS clients

The products listed in the following table can be used for NFS clients.

Table 13-1 Products supported for NFS Clients

Product name	NFSv2 or NFSv3 client	NFSv4 client ^{#1}	IPv6 connectivity
AIX 5L V5.2	Y	--	--
AIX 5L V5.3 (5300-09 or later)	Y	Y	--
AIX V6.1	Y	Y	Y
AIX V7.1	Y	Y	Y
AIX V7.2	Y	Y	--
CentOS 5.4	Y	--	--
CentOS 6.2	Y	--	--
CentOS 6.3	Y	--	--
FreeBSD 5.4	Y	--	--
FreeBSD 6.1	Y	--	--
HP-UX 11i	Y	--	--
HP-UX 11i v2 (PA-RISC)	Y	--	--
HP-UX 11i v3 (HP-UX 11i-OE B.11.31 or later, HP-UX 11i-OE.OE B.11.31 or later)	Y	Y	Y
HP Tru64 UNIX 5.1	Y	--	--
IRIX 6.5	Y	--	--
Mac OS X v10.8	Y	--	--
Mac OS X v10.11	Y	--	--
Oracle Direct NFS Client (Oracle Database 11g Release 2 (11.2.0.3.0) for Linux (x86))	Y ^{#2}	--	--
Oracle Direct NFS Client (Oracle Database 11g Release 2 (11.2.0.3.0) for Microsoft Windows (x64))	Y ^{#2}	--	--
Oracle Linux 6.4	Y	--	--
Oracle Linux 6.5	Y	Y	--
Oracle Linux 6.6	Y	Y	--
Red Hat Enterprise Linux Advanced Platform v5.6 (Linux version 2.6.18-238.el5 or later)	Y	Y	--

Product name	NFSv2 or NFSv3 client	NFSv4 client#1	IPv6 connectivity
Red Hat Enterprise Linux AS v3	Y	--	--
Red Hat Enterprise Linux AS v4 (x86)	Y	--	--
Red Hat Enterprise Linux Server v4.5	Y	--	--
Red Hat Enterprise Linux Server v5.1	Y	--	--
Red Hat Enterprise Linux Server v5.3	Y	--	--
Red Hat Enterprise Linux Server v5.4	Y	Y	--
Red Hat Enterprise Linux Server v5.8	Y	--	--
Red Hat Enterprise Linux Server v6.1	Y	Y	--
Red Hat Enterprise Linux Server v6.2	Y	Y	--
Red Hat Enterprise Linux Server v6.3	Y	--	--
Red Hat Enterprise Linux Server v6.4	Y	--	--
Red Hat Enterprise Linux Server v6.5	Y	Y	--
Red Hat Enterprise Linux Server v6.6	Y	Y	--
Red Hat Enterprise Linux Server v6.8	Y	Y	--
Red Hat Enterprise Linux Server v7.1	Y#2	--	--
Red Hat Enterprise Linux Server v7.2	Y#2	Y	--
Red Hat Enterprise Linux Server v7.3	Y#2	Y	--
Red Hat Enterprise Linux Server v7.4	Y#2	Y	--
Red Hat Enterprise Linux Server v7.5	Y#2	Y	--
Red Hat Linux 8.0	Y	--	--
Solaris 9 Operating System (SunOS 5.9) SPARC Platform Edition	Y	--	--
Solaris 10 Operating System (SunOS 5.10) SPARC Platform Edition (Solaris 10 10/08 or later)	Y	Y	Y
Solaris 11.2 Operating System (SunOS 5.11) SPARC Platform Edition	Y	--	--
SUSE Linux 8.0	Y	--	--
SUSE Linux 9.0	Y	--	--
SUSE Linux 10 SP1	Y	--	--
SUSE Linux 11 SP1	Y	--	--
SUSE Linux 11 SP2	Y	--	--
SUSE Linux 11 SP3	Y	--	--
Turbolinux 10 Server	Y	--	--

Product name	NFSv2 or NFSv3 client	NFSv4 client ^{#1}	IPv6 connectivity
Ubuntu 8.04 LTS	Y	Y	--
Ubuntu 10.04 LTS	Y	Y	--
Ubuntu 14.04 LTS	Y	--	--
Ubuntu 15.04	Y	--	--
Ubuntu 16.04 LTS	Y	--	--
VMware ESX 4.0	Y ^{#2}	--	--
VMware ESX 4.1	Y ^{#2}	--	--
VMware ESXi 5.0	Y ^{#2}	--	--
VMware ESXi 5.1	Y ^{#2}	--	--
VMware ESXi 5.5	Y ^{#2}	--	--
VMware ESXi 6.0	Y ^{#2}	--	--
XenServer 6.1	Y ^{#2}	--	--

Legend:

Y: Supported --: Not supported

#1

If you use the NFSv4 protocol, the availability of some of the functions provided by the HDI system may depend on the product used for the NFS client. Check the product documentation to determine whether or not each function is available.

#2

Supported for only the NFSv3 protocol.

These products are also supported when VMware ESX 3 or later is used.

KDC server

To use Kerberos authentication for authenticating users, you must use a UNIX server or an Active Directory domain controller for the KDC server.

UNIX server

The following products are supported when a UNIX server is used for the KDC server:

- o AIX 5L V5.3
- o HP-UX 11i v3
- o Red Hat Enterprise Linux Advanced Platform v5.2
- o Solaris 10 Operating System (SunOS 5.10) SPARC Platform Edition

Active Directory domain controller

The following products are supported when an Active Directory domain controller is used for the KDC server:

- Microsoft(R) Windows Server(R) 2008 Enterprise
- Microsoft(R) Windows Server(R) 2008 Enterprise 32-bit
- Microsoft(R) Windows Server(R) 2008 Standard
- Microsoft(R) Windows Server(R) 2008 Standard 32-bit
- Microsoft(R) Windows Server(R) 2008 R2 Enterprise
- Microsoft(R) Windows Server(R) 2008 R2 Standard

ID mapping server

In an NFSv4 domain configuration, when an external server is used to convert the NFS client user names and group names to UIDs and GIDs, respectively, the external server must be either an LDAP server designed for user authentication or an NIS server.

LDAP server for user authentication

The following products are supported when an LDAP server designed for user authentication is used as the ID mapping server

- OpenLDAP 2.2.23
- Sun Java(TM) System Directory Server 5.2

NIS server

In an HDI system, you can use an Active Directory domain controller as well as a UNIX server for the NIS server.

When a UNIX server is used for the NIS server, there are no restrictions on the product version as long as a product with NIS functions is installed on the server.

The following products are supported when an Active Directory domain controller is used for the ID mapping server:

- Microsoft(R) Windows Server(R) 2008 Enterprise
- Microsoft(R) Windows Server(R) 2008 Enterprise 32-bit

Configure Active Directory. To use the GUI, install the ID management tools for UNIX.

Network configurations

This section describes network configurations for an NFS environment either when only the NFS service is running or when both the CIFS and NFS services are running at the same time.

Network configuration when only the NFS service is running

The following shows an example of a network configuration when only the NFS service is running.




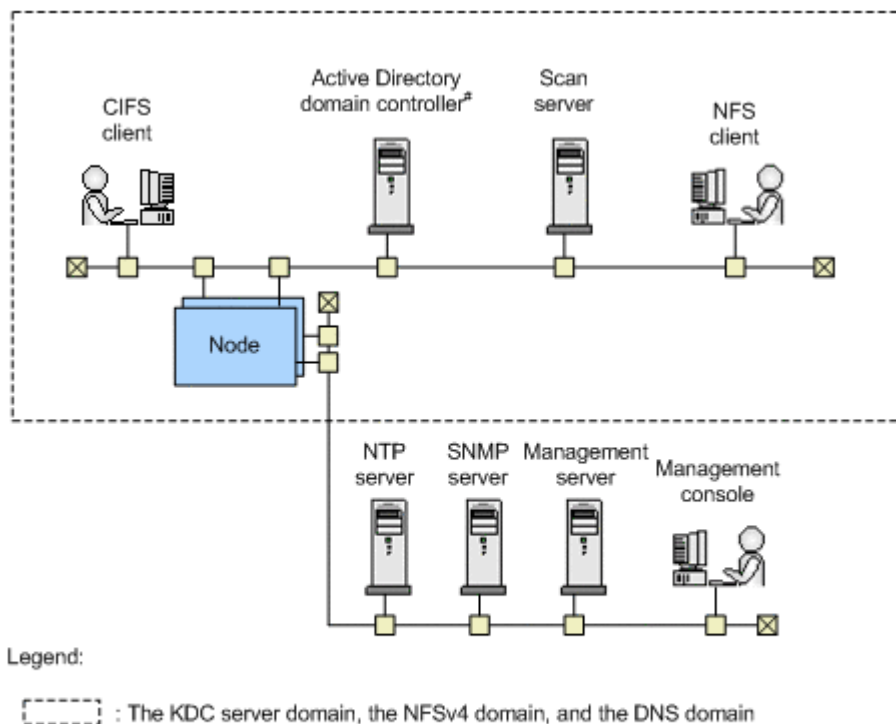
Figure 13-1 Example of a network configuration when only the NFS service is running

You can set up a KDC server domain and an NFSv4 domain for each HDI node. If you use both Kerberos authentication and an NFSv4 domain configuration to run an HDI system, the ranges of KDC server domains, NFSv4 domains, and DNS domains to which the NFS clients belong to must all match.

Network configuration when both the CIFS and NFS services are running

If you run both the CIFS service and the NFS service at the same time and you use Kerberos authentication with the NFS service, you must use an Active Directory domain controller in order to share the KDC server between

the CIFS and the NFS services. If the NFS service uses an NFSv4 domain configuration, you can accommodate the ID mapping server and LDAP server for user authentication in the Active Directory domain controller by using user mapping with the Active Directory schema method, which is also used by the CIFS service. The following shows an example of a network configuration when the CIFS service and the NFS service are both running at the same time and the external server is shared.



#: An Active Directory domain controller can also be used for the DNS server, the KDC server, the NIS server, the LDAP server for user authentication, the user mapping server with the Active Directory schema method, or the ID mapping server.

Figure 13-2 Example of a network configuration when the CIFS and NFS services are both running at the same time and the external server is shared

If you run both the CIFS service and the NFS service at the same time and use both Kerberos authentication and an NFSv4 domain configuration, the ranges for the Active Directory domains, KDC server domains, and NFSv4 domains must all match.

Configuring an NFS environment when Kerberos authentication and an NFSv4 domain configuration are used

This section describes an NFS environment configuration when Kerberos authentication and an NFSv4 domain configuration are used when only the NFS service is running or both the CIFS service and the NFS service are running at the same time.

For the KDC server that is required for Kerberos authentication, which servers are supported depends on what the HDI system is designed to do.

If you are running only the NFS service, you can use either a UNIX server or a domain controller with Active Directory installed for the KDC server. If you are running both CIFS and KDC servers at the same time, you must use a domain controller with Active Directory installed.

This section describes how to configure an NFS environment that is compatible with the HDI system. In the explanation of the configuration procedure, the following terms related to Kerberos authentication are used:

KDC server domain

This consists of a KDC server, the users who will be authenticated by the KDC server, and the servers that use the authentication information. The KDC server domain is also called a realm. When an Active Directory domain controller is used as the KDC server, CIFS clients and NFS clients are authenticated by the KDC server.

Principal

This is a name used to identify a user who will be authenticated by the KDC server. The format of a principal is *user-name@KDC-server-domain-name*.

Keytab file

This is a file containing information about the hosts that will be authenticated by the KDC server. A keytab file created by the KDC server is transferred to the HDI nodes and to the NFS clients.

In advance, you must create a principal for the NFS service as well as a principal for each NFS client user and register them into the keytab file.

Configuring an NFS environment when only the NFS service is running

To configure an NFS environment when only the NFS service is being used:

1. Configure the KDC server and create a keytab file
Configure the KDC server for Kerberos authentication. Also create a keytab file, which is required for Kerberos authentication by the KDC server.
2. Transfer and install the keytab file
Transfer the keytab file created in step 1 to the HDI nodes and the clients.
Merge the contents of the transferred keytab file with the keytab file that is managed by the HDI node. For each client, install the transferred keytab file.
3. Set the service configuration definition and create an NFS share from an HDI node
In the **Access Protocol Configuration** dialog box, specify the settings for Kerberos authentication and the NFSv4 domain. Also, use the **Create**

and Share File System dialog box or the **Add Share** dialog box to create an NFS share.

4. Mount a file system or directory from the NFS client
Mount the file system or directory for which the NFS share has been set in order to make the NFS share accessible.

For details about how to create keytab files and install them on NFS clients, see the applicable product documentation.

Below are detailed descriptions of the above procedure. This procedure assumes that the ID mapping server has already been set up.

Configuring the KDC server and creating a keytab file

To configure the KDC server and create a keytab file:

1. Use a UNIX server or an Active Directory domain controller to configure the KDC server.
2. Create a keytab file at the KDC server.
Platform commands are used in this step. Acquire an Initial Ticket for the `root` user, create necessary principals, and then create a keytab file with an appropriate file name (for example, `/tmp/nfs.keytab`).

Transferring and installing the keytab file

To transfer and install the keytab file on the HDI nodes and the clients:

1. Transfer the keytab file from the UNIX server or the Active Directory domain controller used to create the keytab file to the home directory of the SSH account for HDI (`/home/nasroot`).
To transfer the keytab file from the UNIX server, use the `scp` command. To transfer the keytab file from the Active Directory domain controller, use reliable software that safely copies files.
2. From an HDI node, execute the `nfskeytabadd` command to merge the transferred keytab file with the keytab file already on the HDI node.
The contents of the transferred keytab file are merged with the keytab file that is on the HDI node.
3. Execute the `nfskeytablist` command to check the merged keytab file.
4. Transfer the keytab file from the UNIX server or the Active Directory domain controller that was used to create the keytab file to the appropriate directory on each client (for example, `/tmp`).
To transfer the keytab file from the UNIX server, use the `scp` command. To transfer the keytab file from the Active Directory domain controller, use reliable software that safely copies files.
5. Install the transferred keytab file on each client.

Set the service configuration definition and create an NFS share from an HDI node

To perform this procedure from an HDI node:

1. In the **Access Protocol Configuration** dialog box on the **NFS Service Management** page, specify the necessary settings, such as for Kerberos authentication and for the NFSv4 domain.
Specify the following information:
 - o NFS protocol version supported by the NFS service
 - o Security flavor
 - o Domain name of the NFSv4 domain
 - o KDC server name and KDC server domain name
Specify an IP address or a host name for the KDC server name. You can also specify a host name with a dot (.). If you use an Active Directory domain controller as the KDC server, specify the name of the Active Directory domain controller.
2. Restart the NFS service.
3. In the **Create and Share File System** dialog box or the **Add Share** dialog box, create an NFS share and specify the settings for Kerberos authentication.
For the security flavor that is set on the **NFS** subtab of the **Access Control** tab, you can use the information specified in the NFS service configuration definition as is or specify unique information for the created NFS share.

Mounting from an NFS client

Execute the `mount` command from an NFS client to make the NFS share accessible from the client.

Specify the following options in the `mount` command:

- NFS protocol version used for access (if the client is Solaris, NFSv4 is used by default)
- Security flavor (**sys**, **krb5**, **krb5i**, or **krb5p**)

For details about how to specify options, see the client documentation.

Configuring an NFS environment when the CIFS and NFS services are both running at the same time

The Active Directory domain controller must contain copying software that can safely transfer keytab files.

To configure an NFS environment when the CIFS and NFS services are both running at the same time:

1. Create a keytab file
From the Active Directory domain controller, create a keytab file.

2. Transfer and install the keytab file
Transfer the keytab file created in step 1 to the HDI nodes and to the clients.
From an HDI node, merge the contents of the transferred keytab file with the keytab file on the HDI node. For each client, install the transferred keytab file.
3. Set the service configuration definition and create an NFS share from an HDI node
In the **Access Protocol Configuration** dialog box, specify the settings for Kerberos authentication and the NFSv4 domain. Also, use the **Create and Share File System** dialog box or the **Add Share** dialog box to create an NFS share.
4. Mount a file system or directory from the NFS client
Mount the file system or directory for which the NFS share has been set in order to make the NFS share accessible.

For details on how to create keytab files and install them on NFS clients, see the applicable product documentation.

Below are detailed descriptions of the above procedure. This procedure assumes that an ID mapping server has already been set up.

Creating a keytab file

Use the Active Directory domain controller to create a keytab file.

Acquire an Initial Ticket for the `root` user, create necessary principals, and then create a keytab file with an appropriate file name (for example, `nfs.keytab`).

Transferring and installing the keytab file

To transfer and install the keytab file from the Active Directory domain controller onto the HDI nodes and onto the clients:

1. From the Active Directory domain controller that was used to create the keytab file, transfer the keytab file to the home directory of the SSH account for HDI (`/home/nasroot`) using software that safely copies files.
2. From an HDI node, execute the `nfskeytabadd` command to merge the transferred keytab file.
The transferred keytab file is merged with the keytab file on the HDI node (`/etc/krb5.keytab`).
3. Execute the `nfskeytablist` command to check the merged keytab file.
4. Transfer the keytab file from the Active Directory domain controller to the appropriate directory on each client (for example, `/tmp`) using software that safely transfers files.
5. Install the transferred keytab file on each client.

Set the service configuration definition and create an NFS share from an HDI node

To perform this procedure from an HDI node:

1. In the **Access Protocol Configuration** dialog box on the **NFS Service Management** page, specify the necessary settings, such as for Kerberos authentication and the NFSv4 domain.
Specify the following information:
 - o NFS protocol version supported by the NFS service
 - o Security flavor
 - o Domain name of the NFSv4 domain
 - o KDC server name and KDC server domain name
Specify a server name beginning with an alphabetic character and consisting of alphanumeric characters, hyphens (-), and underscores (_), or specify an IP address. For the KDC server name, specify the name of the Active Directory domain controller.
2. Restart the NFS service.
3. In the **Create and Share File System** dialog box or the **Add Share** dialog box, create an NFS share and specify the settings for Kerberos authentication.
For the security flavor that is set on the **NFS** subtab of the **Access Control** tab, you can use the information specified in the NFS service configuration definition or specify unique information for the created NFS share.

Mounting from an NFS client

Execute the `mount` command from an NFS client to make the NFS share accessible from the client.

Specify the following options in the `mount` command:

- NFS protocol version used for access (if the client is Solaris, NFSv4 is used by default)
- Security flavor (**sys**, **krb5**, **krb5i**, or **krb5p**)

For details about how to specify options, see the client documentation.

Using File Services Manager To Run the NFS Service

The system administrator performs various management tasks in order to properly maintain the HDI system. Within those various management tasks, this chapter describes the operations that require particular attention for the use of the NFS service. The information in this chapter assumes that the File Services Manager GUI is used.

- [File Services Manager setup](#)
- [Configuring network and system information](#)
- [Service configuration definition](#)
- [Managing NFS shares](#)

File Services Manager setup

The system administrator uses File Services Manager to configure the information required to start HDI system operations. The following figure shows the setup procedure for File Services Manager. Among the operations shown in the figure below, this manual focuses on those dealing with the NFS service. For details on the other operations, see the *Administrator's Guide*.




Figure 14-1 File Services Manager configuration procedure

Configuring network and system information

The system administrator can configure the interface information, network information, and linked external server information for each HDI node from the **System Setup Menu** page of the **Network & System Configuration** dialog box.

To use the NFS service, check the following settings:

- NIS server settings
To allow a netgroup to be specified as a host that can access the NFS shares, you must set up an NIS server.

- ID mapping server settings
To use an NFSv4 domain configuration, you must set up an LDAP server for user authentication or an NIS server for the ID mapping server.
- DNS server settings
If you register the host names of the HDI nodes and the NFS client host names (and the KDC server host names in the event that Kerberos authentication is being used) into the DNS server, then name resolution can be centrally performed for the whole domain from the DNS server.

For details on how to specify settings on the **System Setup Menu** page, see the *Administrator's Guide*.

Editing system files directly

The system administrator can directly edit HDI system files from the **Edit System File** page of the **Network & System Configuration** dialog box. For details about how to directly edit these system files and the settings to be specified, see the *Administrator's Guide*.

The following explains the system file to edit and in what case to edit the system file for NFS service use:

/etc/hosts

Edit this system file to use NFS file locking from hosts that access NFS shares.

Service configuration definition

The following table shows what aspects of the NFS service can be managed by a system administrator. For details on service management, see the *Administrator's Guide*.

Table 14-1 Manageable aspects of the NFS service

Service type	Service name	Configuration definition changes	Service maintenance	Start / stop / restart
NFS service	NFS	Y	N	Y

Legend: Y: Can be managed, N: Cannot be managed.

The system administrator must notify the end users (users of the NFS clients) of the settings specified in the NFS service configuration definition, such as the NFS protocol version supported by the NFS service and the security flavor.

The system administrator should be aware of the following before changing the NFS service configuration definition:

- Make sure that the service configuration definitions for both HDI nodes within a cluster are the same.

- If settings such as the NFS protocol version and security flavor will be released or the maximum transfer size will be changed in the NFS service configuration definition, the system administrator must request the administrators of the NFS client hosts to unmount the file system from the NFS clients beforehand. If these settings are changed before the file systems are unmounted, the NFS clients will not be able to access the file system once the NFS service restarts. After changing the configuration definition and restarting the NFS service, the system administrator must request the administrators of the NFS client hosts to mount the file system again.

Changing the NFS service configuration definition

For details about how to change the NFS service configuration definitions and notes on changing the definitions, see the *Administrator's Guide*. This subsection contains additional notes on changing configuration definitions for the NFS service from the **NFS Service Management** page of the **Access Protocol Configuration** dialog box.

Table 14-2 Notes on the NFS service setup from the NFS Service Management page

#	Item	Description and precautions
1	Number of nfsd processes	The number of <code>nfsd</code> processes that will be started changes automatically depending on the system status and the range of specified values.
2	nfsd buffer size	Before changing the maximum transfer size, you must request that the administrators of the NFS client hosts to unmount the file system from the NFS clients. If the UDP protocol is used to perform NFS mounting, the maximum transfer size is limited to 56 KB even when a value of more than 56 is specified.
3	KDC server domain name	If the KDC server is also used as the Active Directory domain controller, the name specified here is also used as the name of the Active Directory domain. If this is different from the name of the Active Directory domain or domain controller used by the CIFS service, the CIFS service must be restarted.
4	KDC server name(s)	If the KDC server is also used as the Active Directory domain controller, the name specified here is also used as the name of the Active Directory domain controller. If this is different from the name of the Active Directory domain or domain controller used by the CIFS service, the CIFS service must be restarted.

Managing NFS shares

This section contains notes on how a system administrator can use File Services Manager to create an NFS share, and how they can modify the NFS share's attributes.

Creating an NFS share and changing the settings

The system administrator can create an NFS share in the **Add Share** dialog box or in the **Create and Share File System** dialog box. For details on how to create an NFS share, see the *Administrator's Guide*. This subsection provides notes about the settings that are specified in the **NFS** subtab of the **Access Control** tab when an NFS share is created. The same information can also be specified in the **NFS** subtab of the **Access Control** tab in the **Edit Share** dialog box.

When you create an NFS share and edit its attributes, you can specify the following information:

- Hosts allowed to access the NFS share
One of the following can be specified for the hosts that are allowed to access the NFS share:
 - Specific host
Specify a host name or IP address.
 - All hosts belonging to a subnetwork or group
Specify the DNS domain name of the DNS domain to which the NFS client belongs, or the IP address of the subnetwork, or the NIS netgroup.
 - All hosts
Specify the wildcard character (*).
- Security flavor for the specified hosts
Select at least one from among **sys**, **krb5**, **krb5i**, or **krb5p** for the permitted authentication method (UNIX (AUTH_SYS) authentication, Kerberos authentication).
To inherit the existing authentication method permitted for a service, select **Use the default settings**.
The security flavor to be used when an NFS client accesses an NFS share depends on the specified `mount` command (`sec` option) or the option's default value when the file system is mounted (NFS mounting) on the NFS client.
- Access permission for the specified hosts
Specify whether both read and write operations are to be permitted or only the read operation is to be permitted for the NFS share.
- Anonymous user mapping
Specify whether to perform anonymous user mapping for all of the users (**For anyone**), for only the root user (**For root user**), or for none of the users (**Not applied**).

- **UID for anonymous mapping** and **GID for anonymous mapping**

Specify the user ID (UID) and group ID (GID) that will be used by an anonymous user.

In an environment where an NFSv4 domain has been set up, an anonymous user is mapped using the UID for **Anonymous user name** and GID for **Anonymous group name**, which are set in the NFS service configuration definition even if **Not applied** is specified. If **For root user** is specified, the UID and GID specified in **UID for anonymous mapping** and **GID for anonymous mapping** are applied only to the `root` user for the result of anonymous user mapping by the NFS service. If **For anyone** is specified, the UID and GID specified in **UID for anonymous mapping** and **GID for anonymous mapping** take precedence over the settings in the NFS service.

Modifying NFS shares

The system administrator can use the **Edit Share** dialog box to modify NFS shares. For details about how to modify NFS shares and notes on modifying shares, see the *Administrator's Guide*. This subsection provides important notes about modifying NFS shares.

- If an item is not modified, the current value is used.

In addition to the above note, see also the notes on creating NFS shares in [Creating an NFS share and changing the settings on page 14-5](#).

Managing NFS Client Users

This chapter explains how to manage NFS client users.

- [User management methods](#)
- [User management when an NFSv4 domain has been set up](#)

User management methods

In an HDI system, you can manage user information for the NFS clients that use the file system by the following means. User names, group names, UIDs, and GIDs make up the user information that is managed.

Table 15-1 NFS client user information management methods

#	Item	Description
1	File Services Manager [#]	Registers user information in order to use File Services Manager to manage the file system users.
2	NIS server	Registers user information in order to use the NIS server to manage the file system users.
3	LDAP server for user authentication	Registers user information in order to use the LDAP server for user authentication to manage the file system users.
4	KDC server	Registers the information necessary for Kerberos authentication. In addition, File Services Manager, an NIS server, or an LDAP server for user authentication must be used to manage user information.

#

Register the same user information in File Services Manager as is in the NFS clients.

User management when an NFSv4 domain has been set up

If you set up an NFSv4 domain, you can limit the NFS clients that can use the NFSv4 protocol for accessing the HDI system to users in the domain.

When a user in an NFSv4 domain accesses the HDI system using the NFSv4 protocol, an ID mapping server or File Services Manager is required in order to convert user names and group names to UIDs and GIDs.

The following figure shows how an NFS share is accessed when an NFSv4 domain has been set up.

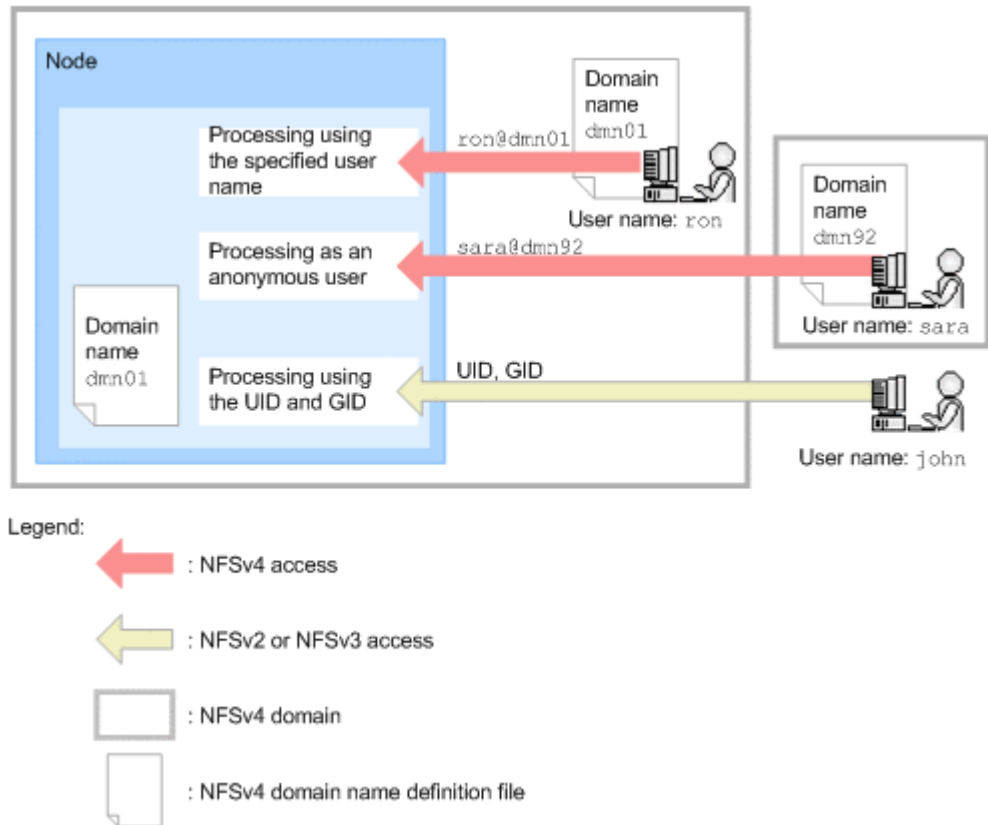


Figure 15-1 Accessing an NFS share when an NFSv4 domain has been set up

An NFSv4 domain consists of HDI node and NFS clients; and only one NFSv4 domain can be set up for each node. If Kerberos authentication is also used, the range of the NFSv4 domain must be the same as the range of the KDC server domain.

To make it so that the NFS clients can participate in the NFSv4 domain to which the node belong, you must configure the NFSv4 domain name from the NFSv4 domain name definition file.

When a user requests access to a node, the user's identification information (*user-name@NFSv4-domain-name*) is used to determine whether they are an NFS client user that is participating in the NFSv4 domain to which the node belongs. When an NFS client user that participates in a different NFSv4 domain issues an access request using the NFSv4 protocol or when ID mapping fails for the user requesting access, access is permitted as an anonymous user. In addition, NFS client users not participating in any NFSv4 domain will be granted access via their UID and GID.

When an NFSv4 domain has been set up, information about the users that access the NFS share is cached temporarily. This cache is retained for 10 minutes. If the cached user information is still available, but the actual user information does not match the cached user information due to a change made to the user information, you must execute the `nfscacheflush` command to access the NFS share.

User Authentication for NFS Clients

This chapter describes how to authenticate NFS client users and contains notes on user authentication.

- [User authentication methods](#)
- [UNIX \(AUTH_SYS\) authentication](#)
- [Kerberos authentication](#)

User authentication methods

The following user authentication methods are supported for the NFS service provided by the HDI system:

- UNIX (AUTH_SYS) authentication
- Kerberos authentication

The system administrator sets the user authentication method and the functions to be used by selecting a security flavor for the NFS service or each NFS share. When a file system for the HDI system is mounted, an NFS client specifies the user authentication method to be used from the security flavor that has been set for the NFS share.

UNIX (AUTH_SYS) authentication

UNIX (AUTH_SYS) authentication is a user authentication method that is executed by the NFS client using the user name and password specified by the user during the login process.

Make sure that users who access file shares by using UNIX authentication do not belong to more than 16 groups. If a user belongs to more than 16 groups, the permissions for the 17th and later groups will become invalid.

Kerberos authentication

This section describes the Kerberos authentication functions supported by the HDI system, as listed below. For the NFS service or each NFS share, you can select one of these functions, as well as UNIX (AUTH_SYS) authentication (`sys`), for the security flavor.

- **krb5**
This is the Kerberos 5 user authentication method.
- **krb5i**
In addition to the Kerberos 5 user authentication method, this selection includes a function that checks the integrity of the data to be transferred.
- **krb5p**
In addition to the Kerberos 5 user authentication method and the use of the function that checks the integrity of the data to be transferred, this selection includes a function for encrypting the data to be transferred.

The security level increases from **krb5**, to **krb5i**, to **krb5p**, but the overhead required also increases. The system administrator must choose the appropriate security flavor by taking into account the operating environment of the HDI system.

When you use Kerberos authentication, it is important that there are no timing errors among the KDC server, HDI nodes, and NFS clients. If a timing error does occur, the NFS clients may not be able to mount a file system or access an NFS share.

Make sure that users who access file shares by using Kerberos authentication do not belong to more than 32 groups. If a user belongs to more than 32 groups, Kerberos authentication for the 33rd and later groups will fail.

Accessing NFS Shares

This chapter describes the procedure for accessing shared directories from an NFS client, and also provides related notes.

- [Access method](#)
- [Mounting and viewing a file system](#)
- [Notes on using a file system from an NFS client](#)

Access method

To access a shared directory from an NFS client, you must mount the file system where the directory resides. For details on how to mount file systems from an NFS client, see [Mounting and viewing a file system on page 17-2](#).

When mounting an HDI file system, it is necessary to specify the host name that corresponds to the virtual IP address of the node.

As such, the host name must be resolvable from both the NFS client and the node, and the virtual IP address obtained from name resolution must be the same on both the NFS client and the node.

To use file locking, specify the host name corresponding to the virtual IP address. If you specify the virtual IP address rather than the host name when mounting the file system, file locking might not function properly.

Mounting and viewing a file system

The file system becomes accessible when a shared directory is mounted from an NFS client. In the case of an NFSv4 client, the root directory can also be mounted, in addition to shared directories.

This section describes how to mount shared directories or the root directory and how the file system appears to NFS clients.

When mounting shared directories

The following figure shows an example of mounting shared directories from the NFS client.

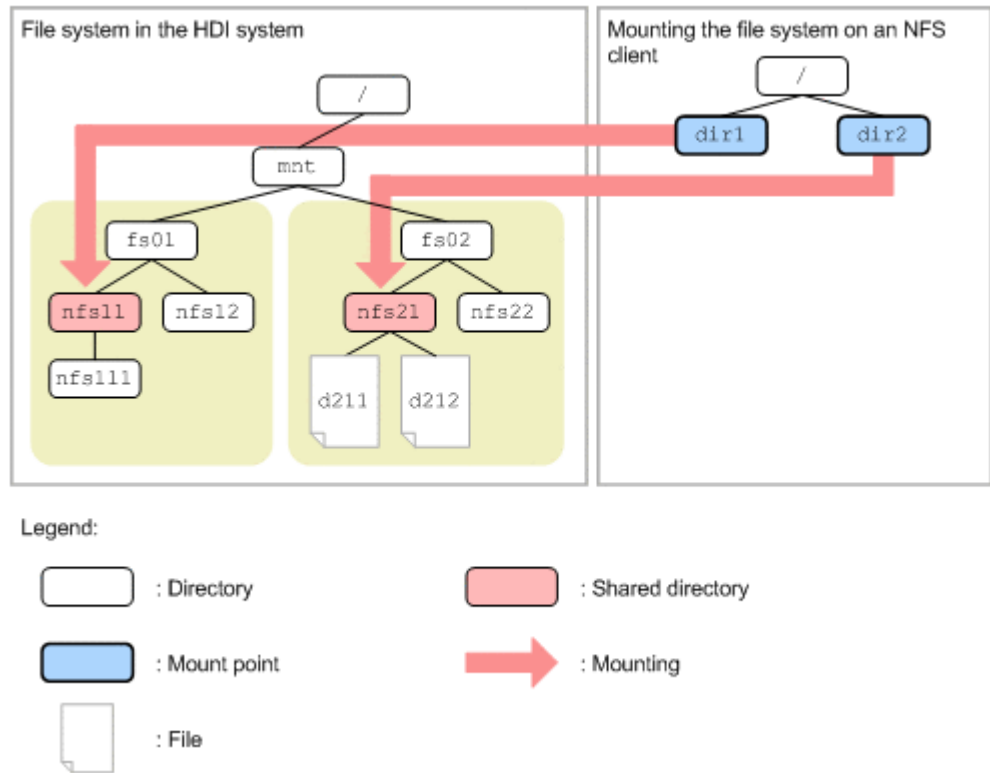


Figure 17-1 Example of mounting shared directories

The following shows an example of executing the `mount` command:

```
mount -o vers=3 node01:/mnt/fs01/nfs11 /dir1
mount -o vers=3 node01:/mnt/fs02/nfs21 /dir2
```

When shared directories are mounted from an NFS client, the directory tree consisting of the directories and files under each shared directory create what is called a file system namespace. To access multiple shared directories, you must mount each shared directory separately.

The following figure shows how the file system appears from an NFS client when shared directories are mounted.

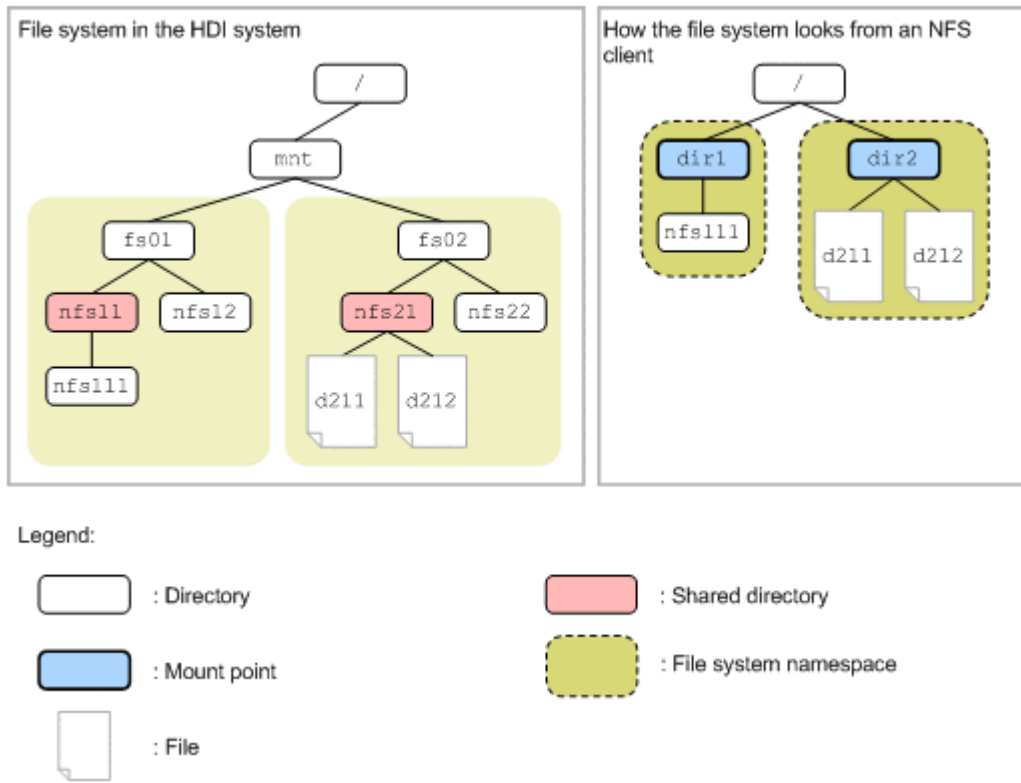


Figure 17-2 How the File System Looks When Shared Directories Are Mounted

When mounting the root directory

When the root directory is mounted from an NFSv4 client, all shared directories under the root directory are also mounted.

The following shows an example of mounting the root directory from an NFSv4 client.

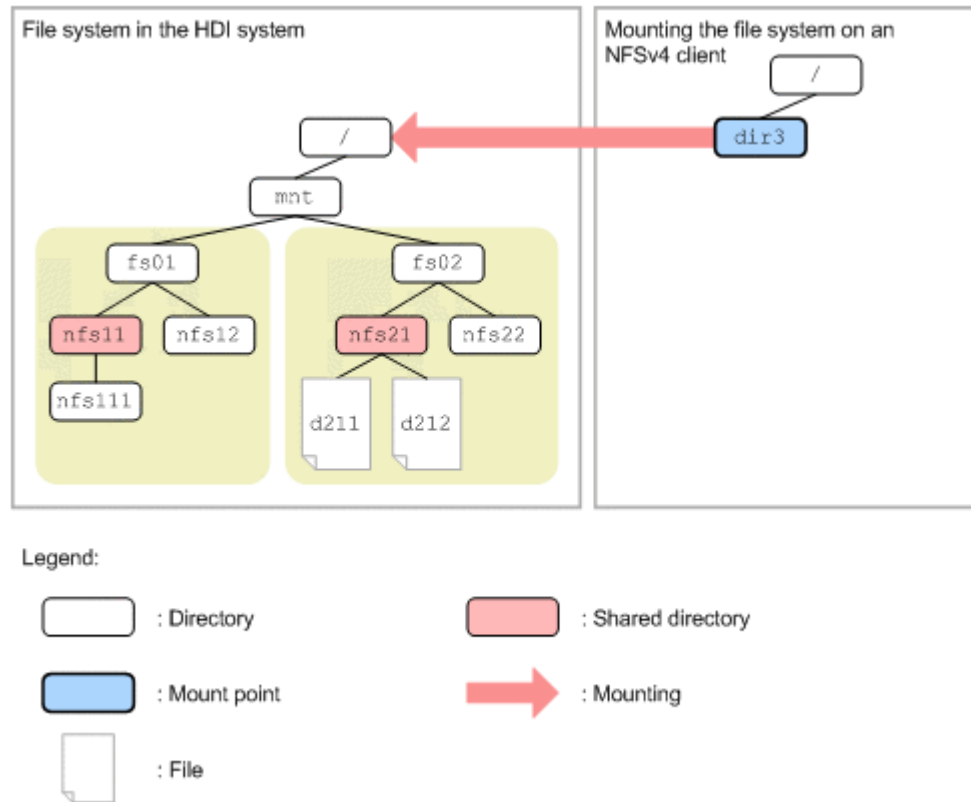


Figure 17-3 Example of mounting the root directory

The following shows an example of executing the `mount` command with the root directory specified:

```
mount -o vers=4 node01:/ /dir3
```

Once the file system's root directory is mounted, NFSv4 clients can access the directory tree as a single virtual file system that consists of multiple NFS shares. Once you mount the root directory, you can access all the shared directories in the same directory tree, thereby eliminating the need to mount individual shared directories.

The NFSv4 client can reference the directories that are located directly between the mount directory and each shared directory, but cannot perform write operations. The files and directories under the directories that are directly connected are hidden from the NFSv4 client.

If the root directory is mounted from an NFS client, the directory tree consisting of the directories and files under all of the shared directories, as well as the directories that are directly between the mount directory and each shared directory will create what is called a file system namespace.

However, if the root directory is mounted from a Solaris 10 or HP-UX 11i v3 NFS client using the NFSv4 protocol, depending on the platform version, directories or files under shared directories might not be displayed as file system namespaces.

The following figure shows how the file system looks from an NFSv4 client when the root directory is mounted.

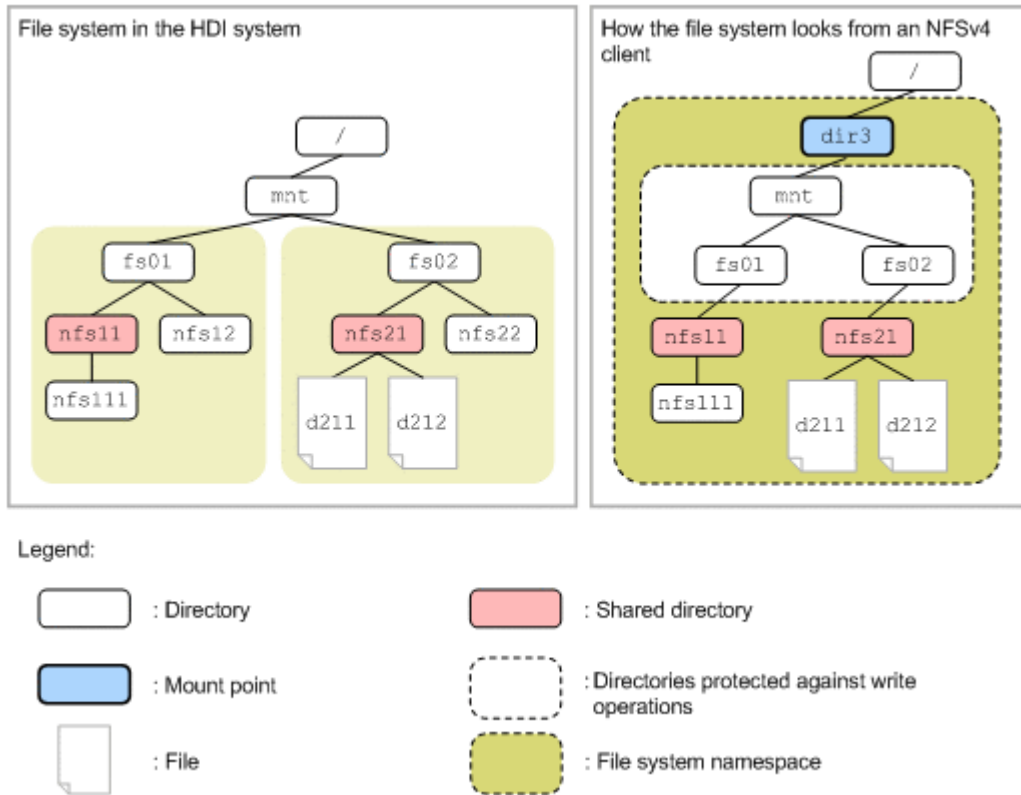


Figure 17-4 How the file system looks when the root directory is mounted

Notes on using a file system from an NFS client

Note the following points when using a file system from an NFS client. For details about notes when modifying HDI settings, see the *Installation and Configuration Guide*.

Also, it sometimes takes some time for an NFS client to access a stub file. Please be aware of this fact when accessing a large amount of files at the same time. For details on stub files, see the *Installation and Configuration Guide*.

Notes on mounting a file system

Note the following points when mounting a file system from an NFS client:

- We recommend that you specify the `hard` option when mounting an HDI system file system from an NFS client. If you specify the `soft` option, and then a failover occurs when an NFS client accesses an HDI system or an NFS client accesses the HDI system that has been failed over, an access request from the NFS client might fail and an `ETIMEDOUT` or `ECONNRESET` error might be output. Note that the `hard` option is specified by default for most NFS clients.

- Explicitly specify the NFSv3 or NFSv4 protocol for the mount option of the NFS client when files 2 GB or greater are used in the file system of an HDI system. If you do not explicitly specify an NFS protocol, the NFSv2 protocol might be automatically used, and files of 2 GB or greater might not be usable.
- Explicitly specify the version of the NFSv4 protocol for the mount option of the NFS client when accessing a file system of the HDI system by using the NFSv4 protocol. If you do not explicitly specify the version of the NFSv4 protocol, depending on the NFS client, the NFSv2 or NFSv3 protocol might be applied.
- You must specify the `hard` and `intr` options when mounting an HDI system file system from an NFS client (using Linux, for example) that cannot interrupt file operations. If you do not specify these options, it might not be possible to stop file operations if an error occurs.
- If you mount an HDI system file system from an NFS client, do not specify a directory that includes non-ASCII characters.
- When deleting or changing a virtual IP address, you must stop all access from NFS clients that are using the virtual IP address, and then unmount the file system on those NFS clients. If you do not do this, the NFS clients will no longer be able to properly access the HDI system.

Notes on using file locking

Note the following points when using file locking from an NFS client:

- To use *file locking* with the NFSv2 or NFSv3 protocol, make sure that resolution of host names and virtual IP addresses of the HDI node and NFS client (in both name-to-address and address-to-name directions) produces a match. To do so, register the host names for virtual IP addresses on the NIS server, the DNS server, or in the `/etc/hosts` files of both the HDI node and the NFS client. Note that after you add, delete, or modify entries in the `/etc/hosts` files, on the NIS server, or on the DNS server, you must restart the NFS services.
- If the host name of the NFS client cannot be resolved from the IP address of the NFS client, lock recovery (file locking processing repeated when a failover occurs or the HDI node is restarted) might not be performed correctly. Ask your vendor whether it is necessary to resolve the host name of the NFS client.
- To use file locking while using the NFSv4 protocol, the NFSv4 domain name specified during the configuration of the NFS service from File Services Manager must be identical to the NFS client NFSv4 domain name.
- If the NFS clients are Mac OS X, when a single file is locked from multiple clients, or a single file is locked with multiple processes from a client that does not support a `SIGLOST` signal, another process that has been waiting to lock the file might acquire the lock when the following processing is executed in an HDI system:
 - Restart of the NFS service, or node OS

- Failover
- When a record lock is collected by using a POSIX lock (segment lock, region lock, or record lock) from an NFS client, an `ENOLCK` error might occur in the following cases:
 - When a subtree check is performed for an NFS share in the HDI system.
 - When all the directories including the parent directory for the lock target file and the NFS mounted directory do not have execution permissions (`x`) for anonymous users in the HDI system.

An `ENOLCK` error does not occur if you specify one of the following settings:

- Among the directories including the parent directory for the lock target file and the NFS mounted directory, add execution permissions (`x`) to the directories that do not have execution permissions (`x`) for anonymous users in the HDI system.
- Use an NFS client to unmount the NFS share directory in the HDI system, and then specify the settings so that a subtree check is not performed for the NFS share. Mount the NFS share directory in the HDI system from the NFS client.
- If an NFS client makes a file-locking request, an `EDEADLK` error will not occur in the following cases. Instead, the job hanging up on the NFS client needs to be canceled.
 - When a deadlock occurs because a Solaris 10 or HP-UX 11i v3 NFS client used the NFSv4 protocol to make a file-locking request
 - When a deadlock occurs because an NFS client used either the NFSv2 or NFSv3 protocol to make a file-locking request for a file that was locked by the NFSv4 protocol
 - When a deadlock occurs because an NFS client used the NFSv4 protocol to make a file-locking request for a file that was locked by either the NFSv2 or NFSv3 protocol
- When the NFS client using Linux locks a file in a directory that was mounted using the TCP protocol, it might take some time to release the wait status for the file lock.

Note that if a process that is locking a file is interrupted from an NFS client running on Linux, the file lock information remains in the HDI system and the system might be unable to lock the corresponding file.

- If an NFS client that is running Linux kernel version 2.4 cancels the process that waits for file locking on the HDI system, the lock information might remain in the HDI system.
- On an NFS client that is running a Linux kernel version 2.4.19 or earlier, the process that waits for file locking might take about 10 seconds to secure the lock.
- If the NFSv2 or NFSv3 protocol is used, the network lock manager (`nlockmgr`) and the network status monitor (`status`) must be running on the NFS client. Execute the `rpcinfo` command as follows on the NFS

client to confirm that `status` and `nlockmgr` are up and running, and listening to the appropriate UDP protocols:

```
rpcinfo -u localhost program-name version
```

If the program is operating properly, `ready` and `waiting` will be displayed. An example is shown below:

```
$ rpcinfo -u localhost nlockmgr 1
program 100021 version 1 ready and waiting
$ rpcinfo -u localhost nlockmgr 3
program 100021 version 3 ready and waiting
$ rpcinfo -u localhost nlockmgr 4
program 100021 version 4 ready and waiting
$ rpcinfo -u localhost status 1
program 100024 version 1 ready and waiting
```

- If an NFS client that is running Linux kernel version 2.4, or 2.6.19 to 2.6.27 locks a file in a directory mounted via the TCP protocol, or if an NFS client that is running a Linux kernel version earlier than 2.4.21 locks a file in a directory mounted via the UDP protocol, the following circumstances will cause the file lock to be released.

- A failover or failback.
- The cluster is stopped and then restarted.

When a file lock is released, other processes are now able to lock the file that was already locked. As a result, the file might become damaged.

Note that when an NFS client that is running Linux kernel version 2.6.19 to 2.6.27 is used, you can prevent the file lock from being released by starting an NFS service from an NFS client, and then mounting the directory via the TCP protocol.

- Depending on the implementation of the NFS client host, when all the following conditions are satisfied, even if a file lock is specified for the data range to be overwritten, contents written in the region before the range might be replaced with 0. You can avoid this problem by setting a file lock in units of transmission length when writing data.
 - You simultaneously write data shorter than the transmission length (`wsize` option at mounting) from multiple clients to a single file.
 - Data is written in the region after the file size and data is written to the same block (the block as seen when the transmission length is used as the unit of length).

Example:

Assume that you mount an HDI file system from NFS client host `x` and NFS client host `y` with a transmission length of 32 KB (with `wsize=32768` and `rsize=32768` specified as options for the `mount` command). Process `A` on NFS client host `x` then places a file lock on bytes 0 to 1,023 of a certain file and writes data to this range. Also, process `B` on NFS client host `y` places a file lock on bytes 1,024 to 2,047 of the same file and writes data to this range.

If process `A` and process `B` run at the same time, the data (contents of bytes 0 to 1,023) written by process `A`, which is closer to the beginning of the file, might be replaced by 0.

Notes on using a file system

Note the following points when using a file system from an NFS client:

- If a failover occurs in an HDI system while an operation (such as creating, updating, or deleting a file or directory of the HDI system) is being performed due to an action (such as a system call, library function, or command) from an NFS client, even if the operation finishes normally in the HDI system, an error might occur on the NFS client.
- When you use the TCP protocol to mount an HDI system file system from an NFS client, access to the file system might require 1 to 10 seconds, depending on the implementation of the host. This situation arises when it has been awhile since the subdirectories and files in the mounted directory have been accessed. Note that although an `ECONNRESET` error might be output to a system log, programs that access the file system using an NFS service will operate normally.
- When you create a special file for the HDI system using an NFS client, note the following points:
 - If you are creating a special file on a file system in an HDI system, the maximum value you can specify for the major number is 4,095, and the maximum value you can specify for the minor number is 1,048,575.
 - If you use the NFSv2 protocol and use an OS other than Linux for an NFS client, and then create a special file using an NFS client, a special file with the major number and minor number different from the specified values might be created. Even when Linux is used, the same problem might occur depending on the distribution. Do not create a special file from an NFS client under the above conditions.
- The NFSv2 protocol is not available for files systems that handle 64-bit inodes. Before setting a file system to handle 64-bit inodes, make sure that no clients are using the NFSv2 protocol for the file system.
- Some applications that run in NFS environments might not support 64-bit inodes. Be sure to configure support for 64-bit inodes only when applications that support 64-bit inodes are used.
- If some or all of the NFS client have multiple network interfaces that are used to communicate with the HDI system, accesses to the NFS share might not be allowed, resulting in an error (`ESTALE` error). Depending on how cluster management software and other software are used, the problem is usually caused by a discrepancy between the IP address used to mount the NFS share and the IP address used to access the NFS share. When you use an HDI system file system from an NFS client in an environment of this type, use one of the following methods to specify where the NFS share is to be published:
 - Use a wildcard (*).
 - Specify the IP addresses of all the network interfaces used on the NFS client.
 - Specify the host names corresponding to all the network interfaces used on the NFS client.

- Specify the IP network that contains the IP addresses of all the network interfaces used on the NFS client.
- Specify the net group that contains the host names corresponding to all the network interfaces used on the NFS client.
- Specify the DNS domain that contains the host names corresponding to all the network interfaces used on the NFS client.
- When an NFS client that is using Solaris accesses a file system for which one of the following types of processing is being executed, a lot of messages might be output to the NFS client environment:

- The file system is being expanded
- A file share is being expanded

You must be careful when setting the rotation (number of files, file size, and so on) for an NFS client system log file.

- When the NFS client using IRIX creates a new directory in the mounted directory and then creates a new file in the new directory, and the same operation is repeated under the new directory, the directories will become deeply nested and the operation might stop.
- When the operation is stopped and the NFS client using HP-UX is copying a file by using the `cp` command, the copy-destination file privilege becomes 000.

When failover occurs in the HDI system and the NFS client using HP-UX updates the file system, the process updating the file system might end abnormally after the failover is performed. To prevent these failures, install PHNE_28568 (for 11.11) from the HP-UX patch programs that is available from the HP (Hewlett-Packard) website. For details on the patch programs, contact the vendor.

- When using Linux kernel for an NFS client host, make sure that you apply the latest patch. If you use the kernel without the latest patch to perform an NFS access, the following problems might occur:
 - An error (error number 528) occurs.
 - Information that is different from the file contents is displayed on the client.
 - Information that is different from the contents written by the client is saved in an HDI system file.
- When using Linux kernel for an NFS client host, an `EBUSY` error might occur while reading a file in an NFS share. If this error occurs, try to access the file again.
- When AIX is used as an NFS client, NFS communication with an HDI system uses a non-privileged port (internet port 1024 or higher) by default. To improve security, we recommend that you use a privileged port (internet port lower than 1024) that only the superuser can create. When privileged ports are used, general users must use an AIX NFS client system to access an HDI system. This improves security when NFS services are used. For details on the settings required for using privileged ports, contact the vendor.

When using a privileged port, configure the settings, such as when creating an NFS share or modifying its attributes, so that the sending port is not restricted. When an NFS share is created using the GUI, the sending port will automatically be setup as a non-restricted port. The system administrator does not need to worry about this. When using commands, specify `do_not_perform` (default) for the `-t` option before executing the `nfscreate` or `nfsedit` command.

- When a computer that is using HP-UX or the RPC program number 100020 (the host on which `rpcinfo -p` displays 100020 in program) is used as an NFS client, you might not be able to correctly reference the contents of the hard link file that is under the directory mounted by an NFS client.

When you create an NFS share or edit NFS share information, and specify the following settings, you can reference the contents of the hard link file correctly. Note that if AIX, IRIX, Linux, or Solaris is used as an NFS client, you do not need to specify the following settings:

- When an NFS share is created
When an NFS share is created using the GUI, the system administrator does not need to worry about the settings.
When using commands, specify `do_not_perform` (default) for the `-s` option of the `nfscreate` command.
- When NFS share information is edited
If the NFS share was created using the GUI, the system administrator does not need to worry about the settings. When using commands, specify `do_not_perform` for the `-s` option of the `nfsedit` command.
- When an NFS access from a client that is running Solaris 10 hangs up, use the driver configuration parameter of Solaris 10 to check the SACK permitted option. If the setting allows using the SACK permitted option (1 or 2 is specified for the `tcp_sack_permitted` parameter in the `ndd` command), an NFS access might hang up. Therefore, specify the setting so that the SACK permitted option is disabled (0 is specified for the `tcp_sack_permitted` parameter).
- When an NFS client accesses a file system of the HDI system, the following message might appear: `file temporarily unavailable on the server, retrying...`. In such a case, the system administrator might be intentionally preventing access to the target file system.
- When the OS on a node is heavily loaded, an NFS client attempting to access an NFS share might generate an `ENOSPC` (no space left on disk) error even before the file system capacity reaches 100%.
- Assume that you run a file update process that, for example, writes updated contents to a temporary file and uses the `mv` command to rename the file (such as the `rsync` command, which is an open source utility), and at the same time another NFS client runs a process for reading the same file. If the processing for reading the file from another NFS client causes contention, the read processing might fail.

- You cannot obtain subtree quota information by executing the command for obtaining quota information from an NFS client. For subtree quota information, contact the system administrator.
- When you use the command for obtaining quota information from an NFS client to obtain quota information for a user who uses HDI file systems, a display overflow might occur if the block usage or quota-related setting exceeds 1 TB.
- When you use the NFSv4 protocol to set an ACL for an Advanced ACL type file system, be sure to permit the `SYNCHRONIZE` mask for any users or groups that might want to do any of the following: update the last access time (`atime`) or last modify time (`mtime`) for a file, or move or rename a file or directory (by using `rename` operation). The reason being, in an Advanced ACL type file system, you need to have the `SYNCHRONIZE` permission for a file in order to update the last access time or last modified time. Furthermore, in order to move or rename a file or directory, you need to have the `SYNCHRONIZE` permission for the following items: the file or directory, the parent directory of the file or directory that will be moved or renamed, and the existing file or directory that will be overwritten.
- When an NFSv4 client using Solaris or HP-UX mounts a subdirectory in a file system that uses the file version restore functionality, that client cannot reference data in the `.history` directory. To reference what was past data at the time of the migration, request that, from the `.history` directory under the applicable subdirectory, the client mount the directory that indicates the date and time at which migration was performed.
- Files or directories used by the system might be displayed on an NFS client. For details about the files and directories used by the system, see [Table 7-3 Notes on the files and folders used by the system on page 7-7](#).
- In a read-write-content-sharing file system, if the same file is updated on multiple HDI systems before the HDI systems are synchronized, a conflict will occur the next time the HDI systems are synchronized. In an HDI system where a conflict occurred, files are saved in the directory in which the files in conflict were originally stored or in the `.conflict` directory just below the mount point of the read-write-content-sharing file system. The end user must check whether the necessary files are saved in the directory that the system administrator told them about. If the files in conflict are not stored in the directory, you must contact the system administrator.

To use the files stored in the `.conflict` directory, copy each file individually to a location of your choosing other than the `.conflict` directory. If the files are copied by directory, incorrect access permissions might be set.

For end users to access the `.conflict` directory, it is necessary to configure settings on the client side so that all files and folders are displayed.

- In the read-write-content-sharing file system, an I/O error might occur if a directory is manipulated remotely from a different site or if a failover, failback, or other temporary failure occurs when a directory is

manipulated. If an I/O error occurs, wait for a while, and then try again. If the error recurs, contact the HDI administrator.

- When the specified MTU value is higher than 1,500, access attempts using NFS from Red Hat Enterprise Linux Server 6.3 or later (including CentOS 6.3 and Oracle Linux 6.4 or later, which are based on Red Hat Enterprise Linux Server 6.3 or later) might cause the system to hang. If this occurs, make sure the kernel configuration parameter `sunrpc.tcp_slot_table_entries` (which is specified on the NFS client to set the multiplicity of the RPC requests) is set to 32. When a value higher than 32 is specified, access attempts using NFS might cause the system to hang.

Files and Directories in an NFS Share

This chapter contains notes about the files and directories that are created in an NFS shared directory.

- [File and directory names](#)
- [ACLs](#)
- [File attributes](#)
- [WORM files](#)
- [Displaying disk capacity](#)

File and directory names

In an HDI system, UTF-8 is used to encode the names of files and directories, which means that the maximum length of a file or directory name is equal to the number of bytes that name takes up when encoded in UTF-8.

The following table describes the maximum length of file and path names on an NFS share.

Table 18-1 Maximum length of file and directory names

No.	Target	Maximum length
1	File name	1,023 bytes
2	Directory name	255 bytes

For an NFS share in a file system that links to an HCP system, the character encoding for the names of files and directories must be set to Unicode (UTF-8).

ACLs

The HDI system supports the following two types of access control lists (ACLs): *Classic ACL* and *Advanced ACL*. Classic ACL can be used to set ACLs that conform to POSIX ACLs. Advanced ACL can be used to set ACLs that conform to Windows NTFS ACLs.

For details about the differences between Classic ACLs and Advanced ACLs, see [Differences between Classic ACLs and Advanced ACLs on page 8-6](#).

In the case of the HDI system, it is recommended that you configure a file system using the Classic ACL type only if the NFS protocol is used for file sharing in the file system; the Advanced ACL type should be used if the CIFS and NFS protocols are both used together or if only the CIFS protocol is used.

NFS client access is controlled by the ACLs and by the access permissions set for the files and directories.

ACLs cannot be referenced or set from NFSv2 or NFSv3 clients. NFSv4 clients can view and set ACLs in the same manner as for CIFS clients.

File attributes

The following table lists the file attributes that can be specified by an HDI system, among the file attributes defined in RFC3530.

Table 18-2 File attributes for the NFSv4 protocol that can be used by an HDI system

File attribute		Available
Mandatory attribute (<i>mandatory</i>)		Yes
Recommended attribute# (<i>recommended</i>)	ACL	Yes
	Creation time	<i>time_create</i>
	DOS file attributes	<i>archive, hidden, system</i>
Named attribute (<i>named</i>)		No

#

In addition to the recommended attributes listed in the above table, there are other recommended attributes that can or cannot be used in an HDI system. Attributes that can be used in an HDI system:

cansetime, case_insensitive, case_preserving, chown_restricted, fileid, files_avail, files_free, files_total, fs_location, homogeneous, maxfilesize, maxlink, maxname, maxread, maxwrite, mode, mounted_on_fileid, no_trunc, numlinks, owner, owner_group, rawdev, space_avail, space_free, space_total, space_used, time_access, time_access_set, time_delta, time_metadata, time_modify, time_modify_set

Attributes that cannot be used in an HDI system:

mimetype, quota_avail_hard, quota_avail_soft, quota_used

WORM files

This section explains NFS share WORM files. Aside from the points listed below, NFS share WORM files have the same features as CIFS share WORM files. For details, see [WORM files on page 8-64](#).

- To make a file read-only in order to create a WORM file, disable the write permission (*w*) for all of the file owners (*user*), the groups the file belongs to (*group*), and all other users and groups (*other*).
- If an attempt is made to apply WORM to a symbolic link file whose link destination file is not a WORM file, the link destination file is changed to a WORM file. Note that the symbolic link file itself is not changed to a WORM file.
- For NFSv2 or NFSv3 clients, you cannot specify a date beyond the year 2038 as the maximum value for a retention period (more precisely, any time after 03:14:07 on January 19, 2038). This limitation is in accordance with the restrictions on these clients. The following lists examples of clients whose maximum retention period is restricted:
 - Distribution using a 32-bit Linux kernel
 - Solaris (32-bit version)

- AIX (32-bit version and versions prior to AIX 5L)
- Client platforms in which the `time_t` type is defined using a signed 32-bit integer (`signed long int`)
- According to the NFS protocol specifications, the `atime` time that can be specified for a file on an NFSv2 or NFSv3 client is a 32-bit unsigned integer. Therefore, the maximum value specifiable as a retention period is February 4, 2106. This rule applies even when a client platform allows a date beyond the year 2038 to be specified as the file's `atime`.
- To delete a WORM file, the read-only attribute must be removed. If a set retention period for a WORM file has expired, the file can be deleted by removing the read-only attribute. However, the data cannot be modified. To remove the read-only attribute, set the write permission (`w`) to one of the following: the file owner (`user`), the group the file belongs to (`group`), or a different user or group (`other`). This makes it impossible to modify the settings of the read permission (`r`) and execute permission (`x`).

Displaying disk capacity

Be careful if the capacity is set to be limited based on hard namespace quotas when an HDI system is linked with an HCP system at the NFS share level. In this case, the hard quota allocated to the namespace is displayed as the disk capacity of the NFS client when the `df` command or `statfs` system call is executed. You can use the `nfsoptset` command to specify whether the hard quota allocated to the namespace is to be displayed as the disk capacity of the NFS client. When the hard quota setting for the namespace is changed, the CIFS client's disk capacity is updated during migrations. Note that if the settings of the hard namespace quotas cannot be obtained or 0 is set as the hard quota, the block capacity and the number of inodes of the file system is displayed as the disk capacity of the CIFS client.

When a block capacity quota is set

When a block capacity quota is set in an HDI system, the disk capacity is displayed on an NFS client, based on the disk usage as shown in the following table.

Table 18-3 Disk capacity displayed on an NFS client when a block capacity quota is set in an HDI system

Quota value Hard limit	Usage	Disk capacity	Disk usage
Not set	--	Block capacity of the file system	Block usage of the file system
Set	At or above the hard limit	Block usage of the namespace	Block usage of the namespace
	Below the hard limit	Hard namespace quota	

Legend:

--: Not applicable

When a quota for number of inodes is set

When a quota for the number of inodes is set in an HDI system, the disk capacity is displayed on an NFS client, based on the disk usage as shown in the following table Note, however, that the number of inodes that has been set cannot be checked from the client.

Table 18-4 Disk capacity displayed on an NFS client when a quota for the number of inodes is set in an HDI system

Quota value Hard limit	Usage	Disk capacity	Disk usage
Not set	--	Total number of inodes in the file system	Number of inodes used in the file system
Set	At or above the hard limit	Sum total of the inodes used in the namespace and the inodes not used in the file system	Number of inodes used in the namespace
	Below the hard limit		

Legend:

--: Not applicable

Notes on Using File Shares

This chapter provides notes on using file systems and file shares that are shared among CIFS, NFS and FTP clients.

- [Notes on accessing file shares](#)
- [Notes on modifying directories](#)
- [Managing users who access file shares](#)
- [Notes on sharing files and directories among CIFS, NFS, and FTP clients](#)

Notes on accessing file shares

Note the following about accessing file shares when files or directories are shared by both the CIFS and NFS services:

- The user IDs (UIDs) and group IDs (GIDs) used by the CIFS service must match the user IDs and group IDs used by the NFS service.

If you use RID or LDAP for user mapping (during the automatic assignment of user IDs and group IDs), first assign user IDs and group IDs that are used by the CIFS service, and then assign those user and group IDs to the applicable user also on the NFS client host.

For user mapping using RID, do not assign the same ID to a CIFS client group ID and an NFS client user ID. Likewise, do not assign the same ID to a CIFS client user ID and an NFS client group ID. If this happens, the CIFS client might not be able to use the CIFS service.

For example, if the range of domain IDs is from 70000 to 100000, the `Domain Users` group ID would be automatically set to 70513. If you then assign the same ID (70513) to an NFS client user ID and access an NFS share from that client, the CIFS clients that belong to `Domain Users` would no longer be able to access the share that was accessed by the NFS client. Similarly, the `Administrator` user ID is automatically set to 70500. If you then assign the same ID (70500) to an NFS client group ID and access an NFS share, the CIFS clients would no longer be able access, as administrators, the share that was accessed by the NFS client. If this happens, you must reassign the user IDs and group IDs for the applicable NFS users, restart the NFS service, and then delete the cached user mapping information from the CIFS service environment.

You can check the user ID and group ID information assigned by user mapping as follows:

For user mapping using RID:

Use the `umapidget` command to check the IDs or names of the users and groups that were mapped by RID.

For user mapping using LDAP (automatic assignment of user IDs and group IDs):

In the **Check for Errors** dialog box, on the **List of RAS Information** page (for `Batch-download`), download the user mapping information.

For details on how to download the user information, see the *Administrator's Guide*.

- A CIFS client cannot access a symbolic link created within the CIFS share. Symbolic links in CIFS shares are created by NFS clients.
- The file and directory permissions set for CIFS shares in the HDI system act in the same way as the access permissions set for NFS shares.
- If multiple NFS and CIFS clients simultaneously attempt to access a file, changes to that file might not be applied. Therefore, set CIFS shares so that they do not use read-only client caches.
- If you restart the NFS service, access from a CIFS client to the file system might fail. In such a case, wait a while, and then attempt to access the file system again.

- If a file for which read-only permissions have been set by the CIFS client is used by the NFS client, the read-only permissions for the file do not take effect on the NFS client.
- If you use non-ASCII characters for file and directory names, you must use Unicode (UTF-8) for the character encoding for the file and directory names used by the NFS clients.
- Because the character code used by an NFS share depends on the NFS client environment, a file or directory name might not be displayed correctly if the NFS client has used other character codes, such as EUC or JIS, to create a file or directory.
- If you use the CIFS service to permit a user who is not a file's owner to change the date and time a file was modified, all users that have write permissions for that file can also change the date and time it was modified via the CIFS client. However, NFS clients only allow file owners to change the date and time a file was modified. Please be aware of this difference when using this setting.
- When a file created by using the NFS service is viewed by a CIFS client, the Linux execution permission (--x) is mapped to the Windows archive attribute. If the file owner's execution permission is deleted from NFS, the Windows system might incorrectly assume that a backup has been completed. For details, see [ACL set by default if there is no inherited ACL on page 8-38](#).
- From an NFS client, do not create files or directories whose names only differ in case from other files or directories already in the same directory. File and directory names are case-sensitive on NFS clients, but not on CIFS clients. As a result, if the only difference between multiple file or directory names is the case, you might not be able to access the intended file or directory from CIFS clients.

Notes on modifying directories

If a CIFS share has been created above an NFS share in the same directory tree, a CIFS client cannot delete or rename the directories between the NFS share and the CIFS shares or the directory for which the NFS share has been created. The following figure shows an example of such a directory tree.

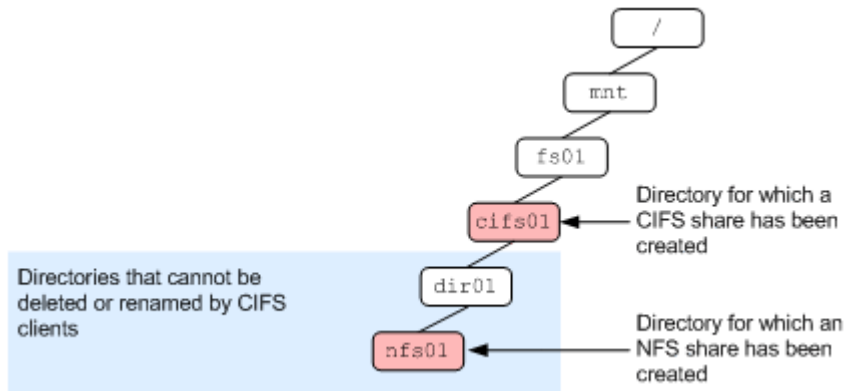


Figure 19-1 Example of a directory tree in which a CIFS share has been created above an NFS share

Managing users who access file shares

User mapping with the Active Directory schema method enables you to manage accounts for the CIFS and NFS services as a single user when files or directories are shared by those services.

The procedure for using user mapping with Active Directory schema method is as follows:

1. Register, into the domain controller, the user ID and group ID of the user who accesses the CIFS share.

For the registration procedure, see the [How to register IDs with Active Directory on page 4-17](#). If Windows Server 2008 or later is used as the domain controller, the `gidNumber` of the user can be used as the group ID instead of the `gidNumber` of the group that the user belongs to.

2. Use the management console to configure settings for making the HDI system join the Active Directory domain.

For details on how to make the HDI system join the Active Directory domain, see the *Administrator's Guide*. Note that, if you registered the `gidNumber` of the user as the group ID in Step 1, you need to use the `cifsoptset` command with the `use_gidnumber` option to change the CIFS service configuration definition to allow the `gidNumber` of the user to be used as the group ID.



Tip: You can check the `gidNumber` of the user in **Primary group name/GID** of the **UNIX Attributes** tab in the Properties page for the Active Directory user.

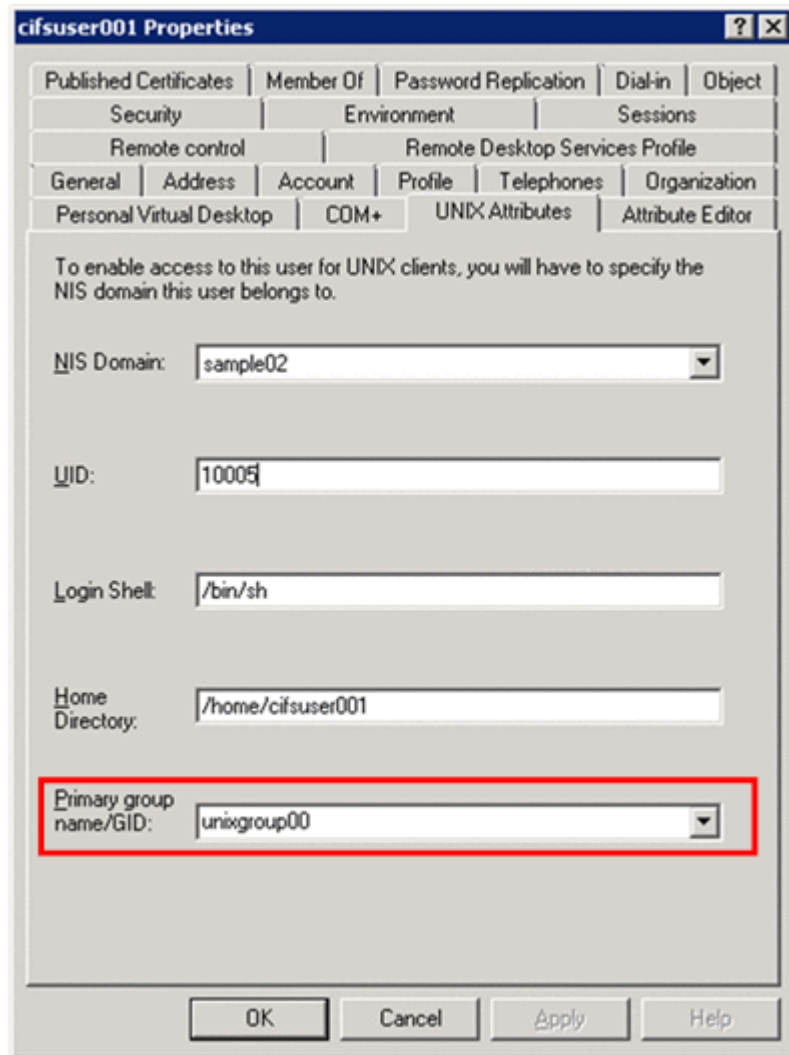


Figure 19-2 Below is an example of a window displaying the UNIX Attributes tab on the Properties page of the user

Notes on sharing files and directories among CIFS, NFS, and FTP clients

The following are notes on sharing files and directories among CIFS, NFS, and FTP clients.

- For NFS and FTP clients, you can create files and directories whose path lengths exceed the maximum length of paths that can be accessed by CIFS clients. If a file or directory is to be shared with CIFS clients, make sure the length of the path including the file name, directory name, host name (or IP address), and share name does not exceed the maximum length of paths that can be accessed by CIFS clients. For details about the maximum length of a file or directory path name, see subsection [Supported characters on page 8-2](#).

- For paths that were created by CIFS clients and include multi-byte characters such as kanji, if the length of such a path exceeds 256 characters (1,024 bytes), attempts to create a symbolic file by specifying that path from an NFS client will result in an error, because the length of that path exceeds the maximum length of paths that can be accessed by NFS clients. If a file or directory is to be shared with NFS clients, make sure the length of the path does not exceed 256 characters (1,024 bytes).

Troubleshooting when using the CIFS service

Detailed information about errors and other items for the CIFS service is output to syslog or the CIFS logs. This section explains the messages output to these logs and the action to be taken for each message.

This section also explains errors and the actions to be taken associated with MMC. This section also contains an FAQ for setting up the CIFS service and CIFS file shares.

- [syslog](#)
- [CIFS logs](#)
- [MMC operation errors and corrective actions](#)
- [File operation errors and corrective actions](#)
- [FAQ](#)

syslog

This subsection explains the messages output to `/var/log/syslog` and the action to be taken for each message.

```
msg=rc0=[error-code] (hosts={ [external-authentication-server-name] }  
[error-details])
```

An error occurred in accessing the external authentication server.

Action:

Check whether the external authentication server has been set up and/or started correctly.

```
[[CHN-number]] error : unable to join. errno:[error-details]
```

When a resource group was started after the CHA name was changed, the resource group was unable to rejoin the NT domain or Active Directory domain.

Action:

Make sure that the domain controller for the NT domain or Active Directory domain can be connected correctly. After checking the connection, restart the CIFS service (for NT domain authentication) or rejoin the domain on the **CIFS Service Maintenance** page (for Active Directory authentication).

```
winbindd environment error. rtn=[error-code]
```

When the CIFS server was started along with a resource group, the trust relationship information could not be obtained by using the RID method of user mapping.

Action:

Make sure that the domain controller for the NT domain or Active Directory domain can be connected correctly.

```
Server: [external-authentication-server-name], [error-  
details].rtn=[error-code]
```

When the AD user mapping is used, an error occurred during a consistency check of the schema method with an external authentication server.

Action:

Check if the name service switch used for the CIFS service user mappings is appropriately set and if the name service switch used by the external authentication server is correct.

```
cifs.init [CHN-number]: Warning. Virtual IP address is not defined.
```

The CIFS service was started (or restarted) without setting a virtual IP address. The CIFS service will start but CIFS access will be unavailable.

Action:

Set a virtual IP address to enable CIFS access.

CIFS logs

This subsection explains the messages output to the CIFS logs (`log.smbd` and `log.winbindd`) and the action to be taken for each message.

log.smbd

This subsection explains the messages output to `/var/log/cifs/log.smbd` and the action to be taken for each message.

```
Failed to join domain: Invalid configuration ("realm" set to
'[specified-domain-name]', should be '[domain-name-registered-on-the-
domain-controller]') and configuration modification was not requested
```

The specified domain name (DNS name) differs from the domain name registered on the domain controller.

Action:

Check the domain controller settings and the domain name (DNS name) that was specified. If the domain controller settings are incorrect, correct them and then retry the operation. If the specified domain name is incorrect, retry the operation with the correct domain name specified.

```
Connection denied from [client-IP-address]
```

Connection from the indicated client was denied.

Action:

Check the following, and either change the setting as required or investigate the status of access from CIFS clients:

- Whether access from the relevant client is denied in **Host access restrictions** or **Host/network based access restriction**.
- Whether the number of connected clients has reached the maximum.

```
allowable_number_of_smbd_processes: number of processes ([number-of-
processes-you-attempted-to-start]) is over allowed limit ([maximum-
number-of-processes])
```

The number of connected clients exceeds the maximum.

Action:

Investigate the status of access from CIFS clients.

```
write_socket_data: write failure. Error = Connection reset by peer
write_socket: Error writing {size-of-writing-data} bytes to socket
{descriptor}: ERRNO = Connection reset by peer
Error writing {size-of-writing-data} bytes to client. {return-
value}. (Connection reset by peer)
getpeername failed. Error was Transport endpoint is not connected
```

The connection from a client was closed.

Action:

The disconnection from the client is due to a timeout or a similar cause. Wait a while, and then retry CIFS access.

```
Failed to verify incoming ticket with error  
smb2: Failed to verify incoming ticket with error
```

User authentication failed in the Active Directory domain.

Action:

Check whether the clocks of the domain controller, HDI node, and CIFS clients match. If these clocks do not match, adjust them so that they do. In addition, before you rejoin the HDI node to the domain, check if CIFS shares are being accessed. If CIFS shares are being accessed, log off and log back on again on the CIFS client.

If CIFS shares are not being accessed, rejoin the HDI node to the domain.

```
Username [user-name] is invalid on this system  
smb2: Username [user-name] is invalid on this system
```

The user account has not been registered.

Action:

Check the following, and change the setting as necessary:

- If you do not use user mapping, the user accounts that are registered in Active Directory (or a CIFS client) must be created in File Services Manager. Check whether the user accounts registered in both File Services Manager and Active Directory (or a CIFS client) are the same.
- If you are using user mapping, possibilities are that the user ID or group ID might not be within the valid range, or that an error might have occurred in accessing the LDAP server. Review the settings related to user mapping. If you are using the Active Directory schema for user mapping, the user ID or group ID might not be registered with Active Directory. Register the required user ID or group ID with Active Directory. For details, see section [How to register IDs with Active Directory on page 4-17](#).

```
create_canon_ace_lists: Some ACEs were skipped. file = [file-path-name], SID = [SID-of-the-relevant-ACE]
```

ACEs whose SIDs could not be converted to UIDs or GIDs were skipped when an ACL was set.

Action:

Check the following, and set the ACL again as necessary:

- This message is output when the ACL includes ACEs for accounts that have been deleted from the domain. If an account does not exist in the domain, an ACE cannot be set because the SID cannot be converted to a UID or GID. Note that when this message is output, all ACEs other than the skipped ACEs are set.

- If this message is output when the accounts exist in the domain, the user mapping functionality might not be functioning correctly. Check the messages output to log.winbindd.

```
create_canon_ace_lists: Can't set ACL. All ACEs were skipped. file =
[file-path-name], SID = [SID-of-the-relevant-ACE]
```

An ACL could not be set because conversion from SIDs to UIDs or GIDs failed for all entries.

Action:

Check the following, and set the ACL again as necessary:

- The user mapping functionality might not be functioning correctly. Check the messages output to log.winbindd.
- This message is also output when all ACEs in the ACL are for accounts that has been deleted from the domain. These ACEs cannot be set because an SID for a non-existent account in the domain cannot be converted to a UID or GID.

```
ads_secrets_verify_ticket: authentication fails for clock skew too
great.
```

Kerberos authentication failed because the system times of the domain controller, HDI node, and CIFS client are out of sync by 5 minutes or more.

Action:

Check the system times of the domain controller, HDI node, and CIFS client, and correct any discrepancies.

log.winbindd

This subsection explains the messages output to `/var/log/cifs/log.winbindd` and the action to be taken for each message.

```
idmap_rid_sid_to_id: [user-or-group-RID] ([UID-or-GID]: [user-ID-or-
group-ID]) too high for mapping of domain: [domain-name] ([minimum-
in-the-domain] -[maximum-in-the-domain])
```

The user ID or group ID to be allocated by user mapping (RID method) is not in the valid range

Action:

Extend or change the range of user IDs or group IDs in the relevant domain.

```
Did not find domain [domain-name]
```

A domain user who is not set by user mapping is attempting access.

Action:

Add a range for the user IDs and group IDs of the domain to which the relevant user belongs.

Cannot allocate `[UID-or-GID]` above `[maximum-user-ID-or-group-ID]!`

The user ID or group ID allocated by user mapping (LDAP automatic allocation) exceeds the valid range.

Action:

Extend or change the range of user IDs or group IDs in the relevant domain.

A `[UID-or-GID]` (`[UID-value-or-GID-value]`) that is out of available range was used (200 - 2147483147). (Name = `[SID]`)

The user ID or group ID registered with the LDAP server by user mapping (LDAP manual allocation) exceeds the allowable range of 200 to 2147483147.

Action:

Specify the user's UID or group's GID value registered with the LDAP server within the allowable range of 200 to 2147483147.

failed to bind to server `ldap://[LDAP-server-IP-address]:[LDAP-server-port-number]` with `dn="[LDAP-server-administrator-DN]"` Error: `[error-details]`

An attempt to access the LDAP server by using user mapping (LDAP method) failed.

Action:

Check whether the specified **LDAP server name** or **LDAP server port number** value is correct, and whether the LDAP server is operating correctly.

`ads_connect` for domain `[NetBIOS-domain-name]` failed: `[error-details]`

An attempt to connect to the domain controller for the Active Directory domain failed.

Action:

Check whether the specified **DC server name(s)** value is correct, and whether the domain controller is operating correctly.

`rpc_np_trans_done`: return critical error. Error was `[error-details]`

The connection to the domain controller for the Active Directory domain or NT domain was closed.

Action:

Check whether the specified **DC server name(s)**, **PDC server name**, or **BDC server name** value is correct, and whether the domain controller is operating correctly.

`cli_start_connection`: failed to connect to `[computer-name-of-the-domain-controller]<20>` (0.0.0.0)

Name resolution of the domain controller failed.

Action:

Register the computer name of the domain controller using a service such as DNS or `lmhosts` so that name resolution can be performed in the HDI system. For details, see the *Installation and Configuration Guide*.

A `[UID-or-GID] ([UID-value-or-GID-value])` that is out of available range was used (200 - 2147483147). (Name = `[sAMAccountName-attribute-value]`)

The user ID or group ID registered with the external authentication server by AD user mapping exceeds the allowable range of 200 to 2147483147.

Action:

Specify the user's UID or group's GID value registered with the external authentication server within the allowable range of 200 to 2147483147.

Could not get unix ID of SID = `[SID to Convert From]`, name = `[user-name]`, type = 30000000

An attempt to get the user ID of user-name has failed.

Causes and corrective actions:

The possible causes and the action to take are described below.

There are irregularities in the Name service switch settings.

Corrective action:

In the **Name service switch** of the **CIFS Service Management** page (**Setting Type:** `User mapping`), review the settings of the CIFS service configuration definition.

The user ID of **user-name** is not registered manually in the domain controller.

Corrective action:

Register the user ID of **user-name** in the domain controller.

Could not get unix ID of SID = `[SID to Convert From]`, name = `[group-name]`, type = 10000000

An attempt to get the group ID of group-name has failed.

Causes and corrective actions:

The possible causes and the action to take are described below.

There are irregularities in the name service switch settings.

Corrective action:

In the **Name service switch** of the **CIFS Service Management** page (**Setting Type:** `User mapping`), review the settings of the CIFS service configuration definition.

The group ID of **group-name** is not registered manually in the domain controller.

Corrective action:

Register the group ID of **group-name** in the domain controller.

The setting that allows the `gidNumber` of the user to be used as the group ID has not been configured.

Corrective action:

Use the `cifsoplist` command to change the setting of the CIFS service configuration definition to allow the `gidNumber` of the user to be used as the group ID.

```
No gidNumber for [SID to Convert From] !?
```

The attempt to obtain the gidNumber of [SID to Convert From] has failed.

Action:

As the group ID of the user that has been registered to the domain controller, specify the `gidNumber` of the user. Alternatively, use the `cifsoplist` command to change the setting of the CIFS service configuration definition to use the `gidNumber` of the user as the group ID.

```
Could not fetch our SID - did we join?
```

An attempt to join the Active Directory domain has failed.

Action:

On the **CIFS Service Maintenance** page, click the **Rejoin Active Directory Domain** button to rejoin the Active Directory domain.

```
add_failed_connection_entry: added domain domain-name (IP-address)
to failed conn cache
```

The authentication processing from the CIFS client might be delayed because there are domain controllers in the domain that cannot communicate with HDI system.

Action:

Make sure that the domain controller that corresponds to **IP-address** is operating correctly and can communicate with HDI system.

If the processing of connecting from HDI system to the domain controller during authentication from a CIFS client times out frequently, you can reduce delays in the processing of CIFS client authentication by suppressing communication to the target domain controller. Use the `cifsopset` command and specify the option to suppress communication to the domain controller that corresponds to **IP-address**.

MMC operation errors and corrective actions

This section summarizes the errors that can occur during operations on CIFS shares from Microsoft Management Console (MMC). In particular, this section covers errors whose cause might be difficult to determine from Windows error

messages. This section also provides detailed information about the causes and corrective actions.

The screenshots appearing in this section are based on MMC 3.0 in Windows Server 2012 R2.

Errors occurring when a share is added

The following are descriptions of an error that might occur when the adding of a share fails, because access was denied by operations being performed from MMC.

Screenshot of the error:




Figure A-1 Window appearing when creating a share fails

Causes and corrective actions:

Possible causes, the message that appears when the error occurs, and the corrective actions are described below.

No permission for the operation

Explanation:

You do not have the required permission to perform an operation on the CIFS share from MMC.

Message output to `/var/log/syslog`:

None

Message output to /var/log/cifs/log.smbd:
None

Corrective action:

Perform the operation as a CIFS administrator who has been registered with File Services Manager.

Invalid file system specified

Explanation:

You have specified the file system created on the other node as the path to the share.

Message output to /var/log/syslog:

```
cifs_addshare : Invalid filesystem specified  
(filesystem=specified-file-system-name). Filesystem belongs  
to CHN[CHN-number]. Own CHN is CHN[CHN-number]
```

Corrective action:

Specify as the path to the share, the file system created on the node on which the operation will be performed.

Non-existent file system specified

Explanation:

A non-existent file system has been specified as the path to the share.

Message output to /var/log/syslog:

```
cifs_addshare : error /enas/bin/  
cifs_fsname2chnnum(ret=2, fsname=specified-file-system-name)
```

Corrective action:

Specify the correct path to the share.

Abnormal termination of the cifscreate command

Explanation:

The `cifscreate` command used to create the share terminated abnormally.

Message output to /var/log/syslog:

```
cifs_addshare : error cifscreate: [return-value-for-the-  
cifscreate-command].
```

Corrective action:

Eliminate the cause of the `cifscreate` command error. Since the `cifscreate` command log data is output to the File Services Manager trace log (`/enas/log/management.trace`), check the File Services Manager trace log.

Errors occurring when the property of a share is changed

The following are descriptions of an error that might occur when an attempt to change a share's properties fail, because access was denied by operations being performed from MMC.

Screenshot of the error:




Figure A-2 Window appearing when an operation on the property of a share fails

Causes and corrective actions:

Possible causes, the message that appears when the error occurs, and the corrective actions are described below.

No permission for the operation

Explanation:

You do not have the required permission to perform an operation on the CIFS share from MMC.

Message output to `/var/log/syslog`:

None

Message output to `/var/log/cifs/log.smbd`:

None

Corrective action:

Perform the operation as a CIFS administrator who has been registered with File Services Manager.

Abnormal termination of the `cifsedit` command

Explanation:

The `cifsedit` command used to edit the share terminated abnormally.

Message output to `/var/log/syslog`:

```
cifs_chgshare : error cifsedit: [return-value-for-the-cifsedit-command].
```

Corrective action:

Eliminate the cause of the `cifsedit` command error. Since `cifsedit` command log data is output to the File Services Manager trace log (`/enas/log/management.trace`), check the File Services Manager trace log.

Errors occurring when a share is removed

This subsection describes errors that can occur when a share is being removed (deleted) by using MMC.

Stopping a share fails due to access denial

The causes of failures to remove a share due to denied access are summarized here.

Screenshot of the error:




Figure A-3 Window appearing when a share is not removed

Causes and corrective actions:

Possible causes, the message that appears when the error occurs, and the corrective actions are described below.

No permission for the operation

Explanation:

You do not have the required permission to perform an operation on the CIFS share from MMC.

Message output to `/var/log/syslog`:

None

Message output to `/var/log/cifs/log.smbd`:

None

Corrective action:

Perform the operation as a CIFS administrator who has been registered with File Services Manager.

Abnormal termination of the `cifsdelete` command

Explanation:

The `cifsdelete` command used to stop the share ends with an error.

Message output to `/var/log/syslog`:

```
cifs_delshare : error cifsdelete: [return-value-for-the-cifsdelete-command].
```

Corrective action:

Eliminate the cause of the `cifsdelete` command error. Since `cifsdelete` command log data is output to the File Services Manager

trace log (/enas/log/management.trace), check the File Services Manager trace log.

Disconnecting a session fails due to access denial

The causes of failures to close a session due to denied access are summarized here.

Screenshot of the error:




Figure A-4 Window appearing when a session is not closed

Causes and corrective actions:

Possible causes, the message that appears when the error occurs, and the corrective actions are described below.

No permission for the operation

Explanation:

You do not have the required permission to perform an operation on the CIFS share from MMC.

Message output to /var/log/syslog:

None

Message output to /var/log/cifs/log.smbd:

None

Corrective action:

Perform the operation as a CIFS administrator who has been registered with File Services Manager.

Non-existent session

Explanation:

The session you have tried to close does not exist.

Message output to /var/log/syslog:

None

Message output to /var/log/cifs/log.smbd:

None

Corrective action:

Display the latest information to check the status of the session.

Errors that occur when an open file is closed

The following is a description of an error that might occur when an attempt to close an open file from MMC fails.

Screenshot of the error:




Figure A-5 The window that appears when an attempt to close a file fails

Causes and corrective actions:

A possible cause, the message that appears when the error occurs, and the corrective action are described below.

No permission for the operation

Explanation:

You do not have the permission required to perform an operation on the open file from MMC.

Message output to `/var/log/syslog`:

None

Message output to `/var/log/cifs/log.smbd`:

None

Corrective action:

Perform the operation as a CIFS administrator who has been registered with File Services Manager.

Error occurring when a session is displayed

The following is a description of an error that might occur because the remote procedure call failed when a session was displayed from MMC.

Screenshot of the error:




Figure A-6 Window appearing when displaying a session fails

Causes and corrective actions:

A possible cause, the message that appears when the error occurs, and the corrective action are described below.

A timeout occurred

Explanation:

A timeout occurred on the CIFS client because it took too much time to display the session information.

Message output to /var/log/syslog:

None

Message output to /var/log/cifs/log.smbd:

None

Corrective action:

Specify a longer timeout time for the CIFS client.

For details about how to specify a longer timeout time, contact Microsoft Support.

File operation errors and corrective actions

This section summarizes the errors that can occur during operations on CIFS shares from Explorer. In particular, this section covers errors whose cause might be difficult to determine from Windows error messages. This section also provides detailed information about causes and corrective actions.

If the error message "The system cannot find the path specified." is displayed




Figure A-7 Example display of the error message "The system cannot find the path specified."

Causes and corrective actions:

The possible causes and the actions to take are described below.

A communication error occurs between a client and an HDI system

Corrective action:

Check the connection between the client and the HDI system. Attempt to access the HDI system again by reconnecting to the HDI system and logging into the client again.

A network drive on the client side is disconnected

Corrective action:

Check whether the network drive is disconnected, and re-map it if it is.

If the error message "An unexpected network error occurred." is displayed

The following error message might be displayed when the CIFS service is accessed from a network drive allocated to a client.

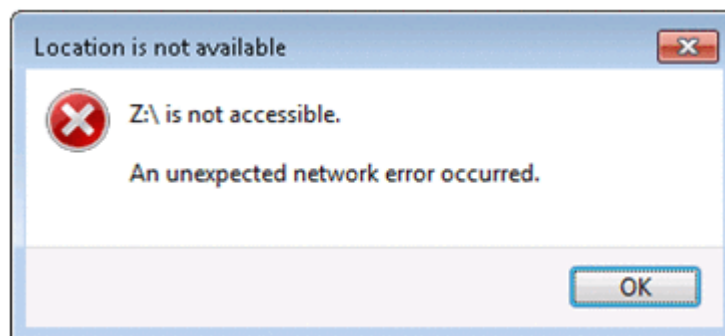


Figure A-8 Example display of the error message "An unexpected network error occurred."

Causes and corrective actions:

The possible causes and the action to take are described below.

The session was temporarily disconnected due to the CIFS service being restarted, or a failover or failback occurred while the CIFS service was being accessed.

Corrective action:

Temporarily deallocate the network drive allocated to the client, re-allocate it, and then access the network drive again.


If the error message "Insufficient system resources exist to complete the requested service." is displayed

Figure A-9 Example display of the error message "Insufficient system resources exist to complete the requested service."

Causes and corrective actions:

The possible causes and the action to take are described below.

System resources at the client have become insufficient because file handles were repeatedly opened and closed over a short period of time.

Corrective action:

Restart the client or, after a little while, retry the operation.

If the error message "A device attached to the system is not functioning." is displayed




Figure A-10 Example display of the error message "A device attached to the system is not functioning."

Causes and corrective actions:

The possible causes and the action to take are described below.

An error might occur if all the following conditions are met:

- The file system is a home-directory-roaming file system.
- The function for automatically creating home directories is disabled.
- When a user attempts to logs in, the home directory of the user does not exist.

Corrective action:

CIFS administrators

The communication error with HCP might have occurred. Make sure that KAQM37526-E is output in the system message, and follow the instructions in the message to fix the issue. After the connection with HCP is restored, log in to the system again.

General users

Use either of the following two methods:

Method 1

Ask the CIFS manager to manually create a home directory for the user.

Method 2

For each linked HDI system, ask the HDI manager to use the `cifsedit` command to enable the function for automatically creating home directories.

If the error message "This server's clock is not synchronized with the primary domain controller's clock." is displayed

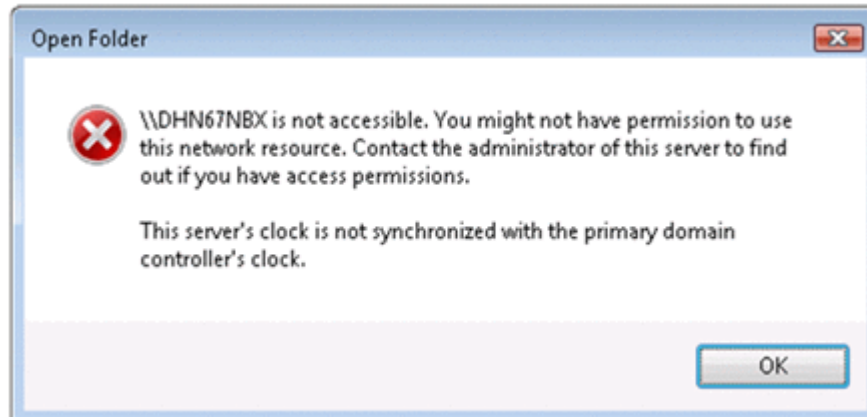


Figure A-11 Example display of the error message "This server's clock is not synchronized with the primary domain controller's clock."

Causes and corrective actions:

The possible causes and the action to take are described below.

User authentication failed because the system times of the domain controller, HDI node, and CIFS client are not synchronized.

Corrective action:

Synchronize the system times of the domain controller, HDI node, and CIFS client. After synchronizing the system times, log off of the CIFS client and then log in again.

A Windows client using SMB 2.0 cannot view folders or files in a CIFS share

Example:

- A folder or file created in a CIFS share cannot be viewed by a client.
- Even if you try deleting a file in a CIFS share, the file never actually gets deleted.

Causes and corrective actions:

The possible causes and the action to take are described below.

This is due to a temporary failure or a problem that occurred in Windows.

Corrective action:

Use either of the following two methods:

Method 1

Wait a while, and then try again.

Method 2

Report to Microsoft Support that a folder or file cannot be viewed using SMB 2.0 and ask for assistance.

If an error message is displayed when a folder or file in a CIFS share is accessed or when the `net view` command is executed to display a list of share names

Example:

- When the `net use` command is executed, the following message is displayed:
"System error 2148073478 has occurred"
- When the `dir` command is executed for a UNC path, the following message is displayed:
"Invalid Signature"
- In cases other than the above, the following error message is displayed:
"An extended error has occurred"
- When the `net view` command is executed, the following message is displayed:
"System error 53 has occurred. The network path was not found."

Causes and corrective actions:

The possible causes and the action to take are described below.

The CIFS client issued a Secure Negotiate request via SMB 3.0 to an HDI system that is configured to not use SMB 3.0.

Corrective action:

On the **CIFS Service Management** page (**Setting Type:** `Basic`) of HDI, configure settings so that all access attempts by the CIFS client use SMB 3.0.

When the following error message is displayed during file creation: The specified server cannot perform the requested operation

Causes and corrective actions:

The possible causes and the action to take are described below.

- The size of the file to be created exceeds the size of the unused capacity of the file system.
- The block usage of the quota exceeds the hard limit when a file of the specified size is created.

Corrective action:

Check the usage of the file system and the block usage of the quota, and then increase the amount of free space in the file system or block. Alternatively, create the file in another file system or folder.

Note that, by using the `cifsoptset` command to specify the settings so that the file system can be checked to see whether it has sufficient free space before data is written to a file to be created or overwritten, a message stating that the disk does not have sufficient free space is sometimes displayed.

A domain user sometimes cannot access a CIFS service and KAQG52019-E or KAQG52018-W system message is output

Causes and corrective actions:

The possible causes and the action to take are described below.

- The times on the domain controller and HDI are not synchronized.
- The SRV records of the DNS server that HDI references include a domain controller to which HDI cannot connect.

Corrective action:

If the times on the domain controller and HDI are not synchronized:

Synchronize the times on the domain controller and HDI.

If the times on the domain controller and HDI are synchronized:

A domain controller that cannot communicate with HDI might exist. Check whether the domain controller displayed in the message (KAQG52019-E or KAQG52018-W) can communicate with HDI. If the domain controller cannot communicate with HDI or is temporarily unable to communicate with HDI, for example, because of a disruption to the communication path or a device failure, resolve the problem so that the domain controller can communicate with HDI.

If you cannot resolve the problem, perform either of the following:

- Remove the domain controller that cannot communicate with HDI from the SRV records of the DNS server.
- Use the HDI routing function to prohibit access to the domain controller that cannot communicate with HDI. This prevents request from being sent to the domain controller. For details on the HDI routing function, see the *CLI Administrator's Guide* or the *Administrator's Guide*.

An attempt to change access permission or ownership of a file or folder (for which access is permitted) results in an error with the "Access is denied" message

Causes and corrective actions:

The possible causes and the action to take are described below.

- The file system has insufficient free capacity.
- The owner or owner group of the file or folder (for which access permission is to be changed) matches one of the following conditions:
 - The block usage of the quota reaches the hardware limit.
 - The block usage of the quota exceeds the grace period that is allowed for when the block usage exceeds the software limit.
- When the owner or owner group of a file or folder is changed, the new owner or owner group matches one of the following conditions:
 - The block usage or inode usage of the quota reaches the hardware limit.

- The block usage or inode usage of the quota exceeds the grace period that is allowed for when the usage exceeds the software limit.

Corrective action:

Check the file system usage and the block usage or inode usage of the quota, and then increase the free capacity. For details about how to check each usage, see the *CLI Administrator's Guide*.

FAQ

This section provides answers to frequently asked questions for setting up the CIFS service and CIFS file shares. For details on how to perform GUI operations used in the procedures described below, see the *Administrator's Guide*.

My system performance sometimes suffers when CIFS file shares are accessed. Is it possible to improve the system performance?

Yes, it is. By default, data is written to the disk drives the moment a close request is issued. However, if a large amount of data is being written to disk drives, your system performance might suffer. You might be able to improve the system performance by changing the settings so that data is written to the disk drives at fixed intervals.

There are two types of settings for handling write requests from CIFS clients: the default setting for all of the file shares and specific settings for individual file shares.

Setting the default setting for all of the file shares so that data is written at fixed intervals:

From **CIFS default setup** in the **CIFS Service Management** page (**Setting Type:** *Performance*), select **Routine disk flush only** for **Disk synchronization policy**.

Setting specific settings for individual file shares so that data is written at fixed intervals:

From the **CIFS** subtab in the **Advanced** tab, which is in the **Edit Share** dialog box, select **Routine disk flush only** for **Disk synchronization policy**. Specifying **Inherit CIFS service default** causes data to be written in accordance with the default settings. When new file shares are added, these same settings can be changed from the **Create and Share File System** dialog box or the **Add Share** dialog box.

For details on how to handle write requests issued from CIFS clients, see the *Administrator's Guide*.

Is there a user account that is similar to a Windows Administrator account? If so, how can I set one up?

Yes, there is. CIFS administrator accounts are similar to Windows Administrator account. On the **CIFS Service Management** page (**Setting Type: Administration**), in the **CIFS service setup** area, users or groups can be registered as CIFS administrators by using the **CIFS administrator name(s)** field. The users registered as CIFS administrators or members of groups registered as CIFS administrators can access all files and folders like Windows Administrator account. These users and group members are treated as root users in HDI system. Note that a user or group merely named Administrator or Administrators is treated as a general user or group regardless of the file system ACL type or user authentication method. Only users and group members registered as CIFS administrators are treated as root users.

If user mapping is being used, specify a domain name with the user name or group name as follows:

domain-name\user-name

@domain-name\group-name

Can only Direct Hosting of SMB be used for the CIFS service?

Yes, it can. You have two choices for how CIFS clients access shares. Among all of the CIFS clients, some of them can use NetBIOS over TCP/IP and the others Direct Hosting of SMB, or all of them can use just Direct Hosting of SMB.

If you want to use only Direct Hosting of SMB for all of the CIFS clients, from **CIFS service setup** in the **CIFS Service Management** page (**Setting Type: Security**), select **Do not use** for **NetBIOS over TCP/IP**.

How can CIFS clients view the Security tab, which allows them to set up or view ACLs?

For file systems that use the Classic ACL type, the **Security** tab in the Properties dialog box of files and folders in file shares might not be displayed from CIFS clients. To display the **Security** tab, select **Yes** for **Enable ACL** in the **CIFS** subtab of the **Access Control** tab in the **Edit Share** dialog box. When new file shares are added, these same settings can be changed from the **Create and Share File System** dialog box or the **Add Share** dialog box.

For file systems that use the Advanced ACL type, the **Security** tab is always displayed.

Can I specify access permissions for entire file systems?

No. You cannot specify file system access permissions for users or groups, but you can specify file share access permissions. Either of the following can be done for file shares:

- Users or groups can be granted read and write permissions.
- Users or groups can be granted read-only permission.

Note that if a file system with file shares has been mounted as read-only, write permissions will not be enabled for any of the users or groups, even if they normally have write permission.

You can change the access permission settings from **Special permitted users/groups** in the **CIFS** subtab of the **Access Control** tab in the **Edit Share** dialog box. When new file shares are added, these same settings can be changed from the **Create and Share File System** dialog box or the **Add Share** dialog box.

Note that if you are using user mapping, you can only use a command to grant permissions.

Sometimes it takes time to access the CIFS shares. What is the potential cause of this problem?

If the HDI system is accessed from multiple domains, the reason it takes time for users to access CIFS shares is most likely that the specified configuration definition of the CIFS service does not match the configuration of the trust relationship of the domain. Change the configuration definition of the CIFS service to fit the domain configuration. For details about how to change the configuration definition of the CIFS service, see the *Administrator's Guide*.

The error message "Cannot access file" was displayed when I attempted to access a file in a CIFS share while the on-access scan function of the scan software was enabled.

When the on-access scan function of the scan software used by the CIFS client side is enabled, the number of files that are opened simultaneously is increased. As a result, the number of open files reaches the limit, causing an error to occur. This error can be avoided by disabling the on-access scan function.

A SID, not the user name or group name, is displayed in the security tab of the properties window of a file or folder. What causes this problem to occur?

HDI does not support one-way trust relationships. The most likely cause of this problem is that you used a domain that is in a one-way trust relationship. Other potential causes are as follows:

The target user or group was deleted from the domain to which the HDI node belongs.

If the user or group was re-registered after being deleted once, the SID that was assigned when it was re-registered is displayed.

The domain in a two-trust relationship with the domain to which the HDI node belongs cannot be accessed.

The domain was deleted or demoted, or the domain controller might not be connected to a power source. Check the status of the domain.

For details about the domain configuration or registered users and groups, check with the domain administrator.

Microsoft office files that were correctly overwritten and saved on a CIFS client are displayed as temporary files (.tmp) on other CIFS clients. What causes this problem to occur?

The most likely cause of this problem is that Offline Files are used on the CIFS client. Offline Files are also called Client Side Caching (CSC).

To display files from other CIFS clients correctly, you need to synchronize the target files by using Sync Center of the CIFS client where the files were saved. For details, contact Microsoft Support.

You can prevent this problem from recurring by disabling Offline Files. Take one of the following measures:

- In HDI, change the CIFS share settings so that the updated data in the CIFS share file is not cached on the client. Note, however, that if you change the settings so that the updated data in the CIFS share file is not cached on the client, access from CIFS clients will take a long time. For details on how to change the settings, see the *Administrator's Guide*.
- Change the CIFS share settings from MMC to make the files and programs in the CIFS share unavailable to offline users. For details on how to change the settings, see section [Changing CIFS share information on page 9-6](#).
- On the CIFS client, disable Offline Files. For details, contact Microsoft Support.

I see the phenomenon that the file I just created gets invisible or the file I just deleted is still visible in the share. What is the cause of these phenomena?

SMB2 Client Redirector Cache is most likely responsible for the visibility of the files what you just operated on the share, as the SMB2 feature caches the file and directory information on the client. You can disable this feature by manipulating registry on the client. Please contact Microsoft Support Service for details.

Troubleshooting when using the NFS service

This section describes errors that might occur while using the NFS service and the required actions to take.

- [Kerberos authentication errors](#)
- [Errors in an NFSv4 domain configuration](#)

Kerberos authentication errors

This subsection explains possible causes of errors that might occur when Kerberos authentication fails and the actions to be taken.

Cause:

A file system failed to mount, causing a Kerberos authentication failure result to be stored in the cache. After that, the user tried to mount the file system again while the cache was still valid.

Corrective action:

Use the `nfscacheflush` command to delete the information stored in the cache. The amount of time the cache for Kerberos authentication results remains valid depends on the amount of time the NFS service tickets issued from the KDC server remain valid. Usually, the NFS service tickets remain valid for 8 to 10 hours.

Cause:

The actual user information and the user information stored in the cache do not match due to changes in the user information. The NFS share was then accessed while the cache was still valid.

Corrective action:

Use the `nfscacheflush` command to delete the information stored in the cache. The cache containing information about users who have accessed NFS shares remains valid for 10 minutes.

Cause:

An NFS service ticket has not been sent.

Corrective action:

From the NFS client, execute the `kinit` command and make sure that an NFS service ticket can be acquired from the KDC server.

Cause:

The NFS service ticket that was sent is expired.

Corrective action:

From the NFS client, execute the `kinit` command to re-acquire an NFS service ticket from the KDC server.

Cause:

The user authentication ticket that was sent has expired.

Corrective action:

From the NFS client, execute the `kinit` command to re-acquire a user authentication ticket from the KDC server.

Cause:

The `gssd` daemon is not running on the NFS client.

Corrective action:

Start the `gssd` daemon. If it is already running, restart the daemon.

Cause:

The Kerberos authentication settings on the NFS client are not valid.

Corrective action:

See the documentation for the NFS client platform, and then take corrective action.

Cause:

The system times of the various components that make up the KDC server domain (the KDC servers, the HDI nodes, and the NFS clients) are not synchronized.

Corrective action:

Synchronize the system times for all of the components that make up the KDC server domain. Kerberos authentication cannot be performed if the system times differ by 5 minutes or more.

Cause:

The Kerberos ticket-processing daemon is not running on the KDC server.

Corrective action:

Start the Kerberos ticket-processing daemon. If it is already running, restart the daemon.

Cause:

At least one host name registered in the DNS server is invalid.

Corrective action:

Check the host names of the HDI nodes, the host names of the KDC servers, and the host names of the NFS clients. If there are any invalid host names, correct the host names on the DNS server, making sure they are valid.

Cause:

The KDC server settings are invalid.

Corrective action:

See the documentation for the platform running on the KDC server and take corrective action.

Cause:

The DES-CBC-CRC Kerberos encryption algorithm is not being used for all of the following: the HDI nodes, the KDC servers, and the NFS clients.

Corrective action:

If the DES-CBC-CRC Kerberos encryption algorithm is not being used for all of the applicable components, modify the settings for the necessary service principals.

Cause:

The NFS service configuration definitions or Kerberos authentication settings for the NFS shares are invalid.

Corrective action:

Check the NFS service and NFS share settings, and then make sure that Kerberos authentication is enabled.

Cause:

The HDI node keytab file is invalid.

Corrective action:

Make sure that the KDC server keytab file has been correctly merged with the HDI node keytab file.

If you still cannot resolve the problems by performing the actions described above, collect the following information, and then send it to maintenance personnel:

- All of the log data
- The following files or information about the KDC server and the NFS client:
 - Kerberos configuration file (`krb5.conf`)
 - Host information file (`hosts`)
 - File where the DNS server IP address is specified (`resolv.conf`)
 - Keytab file
 - System logs
 - Startup process information

Errors in an NFSv4 domain configuration

This subsection explains the possible cause of an error that might occur in an NFSv4 domain configuration and the action to be taken.

Cause:

The NFSv4 domain name on the NFS client is invalid.

Corrective action:

Make sure that the NFSv4 domain name in the NFSv4 domain name definition file on the NFS client and the NFSv4 domain name in the NFS service configuration definition match.

If you still cannot resolve the problem by performing the action described above, collect the following information, and then send it to maintenance personnel:

- All of the log data
- NFSv4 domain name definition file on the NFS client

Configuring an NFS environment for Kerberos authentication

This section explains (through the use of execution examples) how to configure an NFS environment for Kerberos authentication.

- [NFS environment to be configured in this appendix](#)
- [Configuring the KDC server and adding NFS service principals](#)
- [Distributing and retrieving keytab files](#)

NFS environment to be configured in this appendix

For an overview and list of prerequisites, see [Configuring an NFS environment when Kerberos authentication and an NFSv4 domain configuration are used on page 13-7](#).

The following figure shows how the NFS environment is configured for the execution examples.

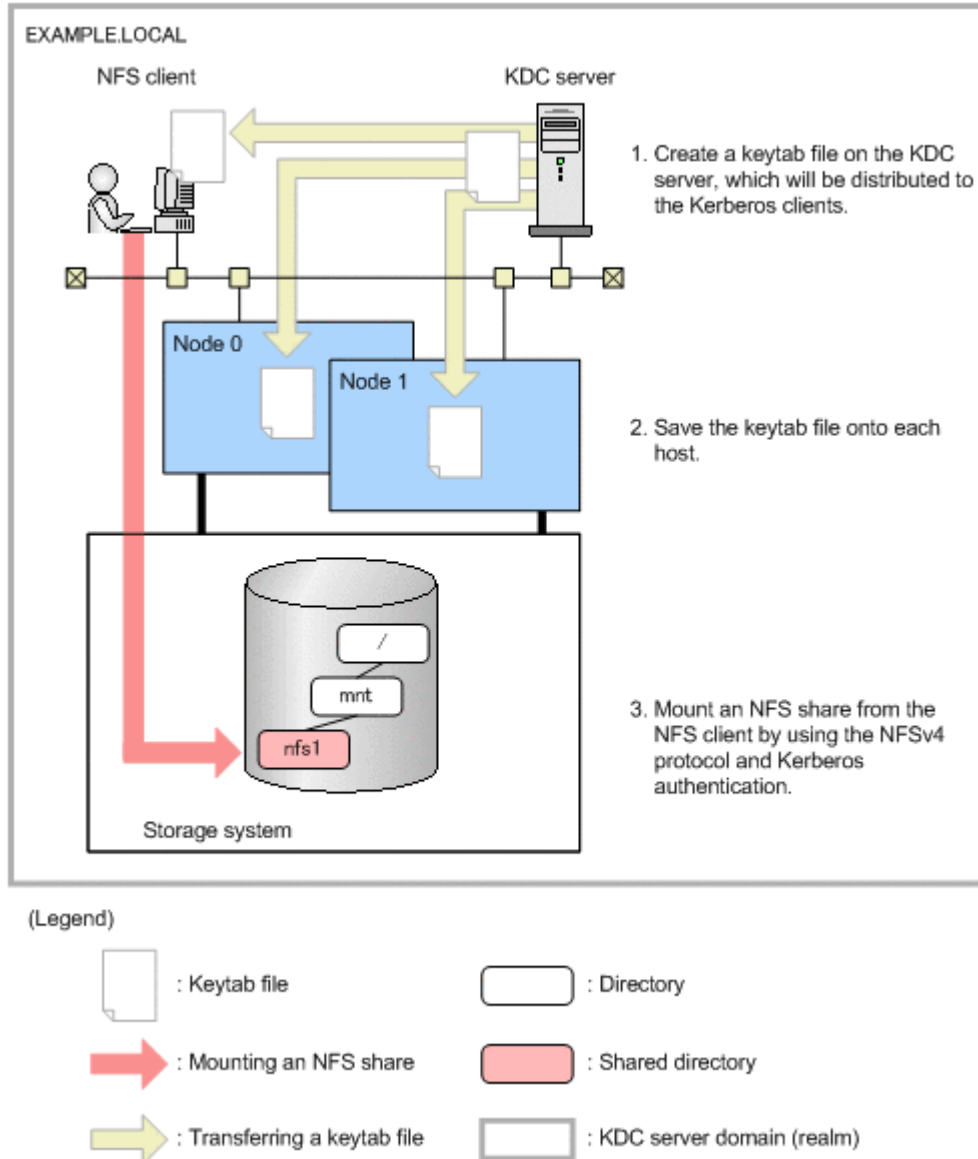


Figure C-1 Example of configuring the NFS environment

The following domain name and keytab file name are used for the hosts in the execution examples.

Table C-1 Domain name and keytab file name used for the hosts

#	Host	Host name (FQDN)	Keytab file name
1	KDC server	kdc1.example.local	..#1
2	Node 0	node0.example.local#2	node0.keytab
3	Node 1	node1.example.local#2	node1.keytab
4	NFS client	cl1.example.local	cl1.keytab

#1

For the execution examples, the keytab files that are distributed to each Kerberos client (host) are created in the `/tmp` directory.

#2

The host name corresponds to the virtual IP address of the HDI node.

Configuring the KDC server and adding NFS service principals

This section describes how to configure the KDC server and add NFS service principals for each platform. The procedures in this section need to be done from the KDC server by an administrator.

Before configuring the KDC server

Before configuring the KDC server, check the following:

- The clocks of all the hosts that are part of the KDC server domain are synchronized.
For Kerberos authentication, if a time difference of 5 minutes or more exists, an error might occur. For Kerberos authentication, we recommend that you use an NTP server.
- The names of all the hosts that are part of the KDC server domain can be resolved via DNS.
All of the host names should be registered with their FQDNs.
- DES-CBC-CRC Kerberos encryption must be used by the HDI nodes, KDC servers, and NFS clients.

For Windows Server 2008

The procedure for configuring a KDC server on Windows Server 2008 and adding NFS service principals is provided below.

To configure the KDC server and add NFS service principals:

1. Configure Active Directory by using the Active Directory wizard.
2. For Windows Server 2008 R2, enable DES encryption.

In Windows Server 2008 R2, because both the DES-CBC-MD5 and DES-CBC-CRC modes of DES encryption are disabled by default, you must enable DES encryption. From **Administrative Tools**, open **Local Security Policy**. Then, from **Security Settings**, select **Local Policies** and **Security Options**, and then double-click **Network security: Configure encryption types allowed for Kerberos**. From the **Local Security Setting** tab, select the **DES_CBC_CRC** check box.

3. Create a user account for mapping NFS service principals, and then add the user account to Active Directory.

Select **Users, New**, and **User** in the Active Directory administrative tool, and then create a user account for each host in the server domain.

This example uses the following user accounts.

Table C-2 Hosts for which user accounts are made and their corresponding user logon names

#	Host	User logon name
1	Node 0	node0
2	Node 1	node1
3	NFS client	c11

In addition, enable DES encryption in the account options for the user accounts. The following figure shows how to specify account options in Windows Server 2008 R2.




Figure C-2 Example of specifying account options (in Windows Server 2008 R2)

4. Execute the `ktpass` command from the command prompt to create a keytab file.

```
> ktpass -princ nfs/node0.example.local@EXAMPLE.LOCAL -mapuser node0 -pass
passwd -crypto DES-CBC-CRC -ptype KRB5_NT_PRINCIPAL -out node0.keytab
> ktpass -princ nfs/node1.example.local@EXAMPLE.LOCAL -mapuser node1 -pass
passwd -crypto DES-CBC-CRC -ptype KRB5_NT_PRINCIPAL -out node1.keytab
> ktpass -princ nfs/cl1.example.local@EXAMPLE.LOCAL -mapuser cl1 -pass
passwd -crypto DES-CBC-CRC -ptype KRB5_NT_PRINCIPAL -out cl1.keytab
```

Use the following `ktpass` command options:

`-princ`

Used to specify an NFS service principal name. (*nfs/host-name-with-FQDN@KDC-server-domain-name*)

`-mapuser`

Used to specify the user name of the account user that was created by using the Active Directory administrative tool.

-pass

Used to specify the password of the account user that was created by using the Active Directory administrative tool.

-crypto

Used to specify the Kerberos encryption algorithm. (You must specify DES-CBC-CRC.)

-ptype

Used to specify the principal type.

-out

Used to specify the name of the keytab file that will be distributed to each Kerberos client (host).

When you execute the `ktpass` command, the account user logon name is mapped to the NFS service principal name. The following figure shows how to perform mapping in Windows Server 2008 R2.




Figure C-3 Example of mapping the user logon name by executing the `ktpass` command (in Windows Server 2008 R2)

For Red Hat Enterprise Linux Advanced Platform v5.2

For this subsection, the following version of Red Hat Enterprise Linux Advanced Platform v5.2 is used:

- Linux version 2.6.18-92.el5 (mockbuild@builder16.centos.org) (gcc version 4.1.2 20071124 (Red Hat 4.1.2-42)) #1 SMP Tue Jun 10 18:49:47 EDT 2008
- Red Hat Enterprise Linux Server release 5 (Tikanga)

To configure a KDC server on a Red Hat Enterprise Linux Advanced Platform v5.2 computer and add NFS service principals:

1. Make sure that the `krb5-server`, `krb5-libs`, and `krb5-workstation` packages have been installed.

```
# rpm -qa | grep krb
krb5-server-1.5-17
krb5-libs-1.5-17
krb5-workstation-1.5-17
```

2. Edit the Kerberos configuration file (`krb5.conf`) so that it matches the file below.

```
# cat /etc/krb5.conf
[logging]
default = FILE:/var/log/krb5libs.log
kdc = FILE:/var/log/krb5kdc.log
admin_server = FILE:/var/log/kadmind.log

[libdefaults]
default_realm = EXAMPLE.LOCAL
dns_lookup_realm = false
dns_lookup_kdc = false
ticket_lifetime = 24h
forwardable = yes

[realms]
EXAMPLE.LOCAL = {
    kdc = kdc1.example.local:88
    admin_server = kdc1.example.local:749
    default_domain = example.local
}

[domain_realm]
.example.local = EXAMPLE.LOCAL
example.local = EXAMPLE.LOCAL

[kdc]
profile = /var/kerberos/krb5kdc/kdc.conf

[appdefaults]
pam = {
    debug = false
    ticket_lifetime = 36000
    renew_lifetime = 36000
    forwardable = true
    krb4_convert = false
}
```

3. Create a KDC database by using the `kdb5_util` utility.

```
# /usr/kerberos/sbin/kdb5_util -r EXAMPLE.LOCAL create -s
Loading random data
Initializing database '/var/kerberos/krb5kdc/principal' for realm
'EXAMPLE.LOCAL',
master key name 'K/M@EXAMPLE.LOCAL'
You will be prompted for the database Master Password.
It is important that you NOT FORGET this password.
Enter KDC database master key:
Re-enter KDC database master key to verify:
```

4. Edit the administrative access control list file (`kadm5.acl`) so that it matches the file below.

```
# cat /var/kerberos/krb5kdc/kadm5.acl
*/admin@EXAMPLE.LOCAL *
```

5. Create an administrative principal.

```
# /usr/kerberos/sbin/kadmin.local -q "addprinc root/admin"
Authenticating as principal root/admin@EXAMPLE.LOCAL with password.
WARNING: no policy specified for root/admin@EXAMPLE.LOCAL; defaulting to
no policy
Enter password for principal "root/admin@EXAMPLE.LOCAL":
Re-enter password for principal "root/admin@EXAMPLE.LOCAL":
Principal "root/admin@EXAMPLE.LOCAL" created.
```

6. Start the Kerberos server daemon.

```
# /usr/kerberos/sbin/krb524d -m
# /usr/kerberos/sbin/krb5kdc
# /usr/kerberos/sbin/kadmind
```

7. Acquire the initial ticket for the administrative principal.
Make sure that the initial ticket could be properly acquired.

```
# kinit root/admin
Password for root/admin@EXAMPLE.LOCAL:
# klist -5
Ticket cache: FILE:/tmp/krb5cc_0
Default principal: root/admin@EXAMPLE.LOCAL

Valid starting    Expires          Service principal
03/26/09 16:08:51 03/27/09 16:08:51  krbtgt/EXAMPLE.LOCAL@EXAMPLE.LOCAL
```

8. Create a KDC server keytab file (krb5.keytab) so that the kadmin utility can be used through network.

```
# /usr/kerberos/sbin/kadmin.local
Authenticating as principal root/admin@EXAMPLE.LOCAL with password.
kadmin.local: ktadd -k /etc/krb5.keytab kadmin/admin kadmin/changepw
Entry for principal kadmin/admin with kvno 3, encryption type Triple DES
cbc mode with HMAC/shal added to keytab WRFILE:/etc/krb5.keytab.
Entry for principal kadmin/admin with kvno 3, encryption type DES cbc mode
with CRC-32 added to keytab WRFILE:/etc/krb5.keytab.
Entry for principal kadmin/changepw with kvno 3, encryption type Triple
DES cbc mode with HMAC/shal added to keytab WRFILE:/etc/krb5.keytab.
Entry for principal kadmin/changepw with kvno 3, encryption type DES cbc
mode with CRC-32 added to keytab WRFILE:/etc/krb5.keytab.
```

9. Create a host principal by using the kadmin utility.

```
kadmin.local: addprinc -randkey host/kdcl.example.local
WARNING: no policy specified for host/kdcl.example.local@EXAMPLE.LOCAL;
defaulting to no policy
Principal "host/kdcl.example.local@EXAMPLE.LOCAL" created.

kadmin.local: ktadd host/kdcl.example.local
Entry for principal host/kdcl.example.local with kvno 3, encryption type
Triple DES cbc mode with HMAC/shal added to keytab WRFILE:/etc/krb5.keytab.
Entry for principal host/kdcl.example.local with kvno 3, encryption type
DES cbc mode with CRC-32 added to keytab WRFILE:/etc/krb5.keytab.

kadmin.local: listprincs
K/M@EXAMPLE.LOCAL
host/kdcl.example.local@EXAMPLE.LOCAL
```

```
kadmin/admin@EXAMPLE.LOCAL
kadmin/changepw@EXAMPLE.LOCAL
kadmin/history@EXAMPLE.LOCAL
kadmin/kdc1@EXAMPLE.LOCAL
krbtgt/EXAMPLE.LOCAL@EXAMPLE.LOCAL
root/admin@EXAMPLE.LOCAL
kadmin.local:
```

10. Use the `kadmin` utility to add the `host` principal to the KDC server keytab file (`krb5.keytab`).

Make sure that the `host` principal has been properly added.

```
kadmin.local: addprinc -randkey host/kdc1.example.local
WARNING: no policy specified for host/kdc1.example.local@EXAMPLE.LOCAL;
defaulting to no policy
Principal "host/kdc1.example.local@EXAMPLE.LOCAL" created.

kadmin.local: ktadd host/kdc1.example.local
Entry for principal host/kdc1.example.local with kvno 3, encryption type
Triple DES cbc mode with HMAC/sha1 added to keytab WRFILE:/etc/krb5.keytab.
Entry for principal host/kdc1.example.local with kvno 3, encryption type
DES cbc mode with CRC-32 added to keytab WRFILE:/etc/krb5.keytab.

kadmin.local: listprincs
K/M@EXAMPLE.LOCAL
host/kdc1.example.local@EXAMPLE.LOCAL
kadmin/admin@EXAMPLE.LOCAL
kadmin/changepw@EXAMPLE.LOCAL
kadmin/history@EXAMPLE.LOCAL
kadmin/kdc1@EXAMPLE.LOCAL
krbtgt/EXAMPLE.LOCAL@EXAMPLE.LOCAL
root/admin@EXAMPLE.LOCAL
kadmin.local:
```

11. Use the `kadmin` utility to create an NFS service principal for each host, and then add the NFS service principal to the keytab file.

```
kadmin.local: addprinc -randkey nfs/node0.example.local
...
kadmin.local: ktadd -e des-cbc-crc:normal -k /tmp/node0.keytab nfs/
node0.example.local
...
kadmin.local: addprinc -randkey nfs/node1.example.local
...
kadmin.local: ktadd -e des-cbc-crc:normal -k /tmp/node1.keytab nfs/
node1.example.local
...
kadmin.local: addprinc -randkey nfs/cl1.example.local
...
kadmin.local: ktadd -e des-cbc-crc:normal -k /tmp/cl1.keytab nfs/
cl1.example.local
...
kadmin.local: quit
```

For Solaris 10

For this subsection, the following version of Solaris 10 is used:

- SunOS 5.10 Generic_137137-09 sun4u sparc SUNW, Sun-Blade-1000

- Solaris 10 10/08 s10s_u6wos_07b SPARC Copyright 2008 Sun Microsystems, Inc. All Rights Reserved. Use is subject to license terms. Assembled 27 October 2008

When configuring a KDC server on a Solaris 10 computer, make sure that DNS has been enabled ahead of time.

To configure a KDC server on a Solaris 10 computer and add NFS service principals:

1. Edit the Kerberos configuration file (`krb5.conf`) so that it matches the file below.

```
# cat /etc/krb5/krb5.conf
[logging]
default = FILE:/var/log/krb5libs.log
kdc = FILE:/var/log/krb5kdc.log
admin_server = FILE:/var/log/kadmind.log

[libdefaults]
default_realm = EXAMPLE.LOCAL
dns_lookup_realm = false
dns_lookup_kdc = false
ticket_lifetime = 24h
forwardable = yes

[realms]
EXAMPLE.LOCAL = {
    kdc = kdc1.example.local:88
    admin_server = kdc1.example.local:749
    default_domain = example.local
}

[domain_realm]
.example.local = EXAMPLE.LOCAL
example.local = EXAMPLE.LOCAL

[appdefaults]
pam = {
    debug = false
    ticket_lifetime = 36000
    renew_lifetime = 36000
    forwardable = true
    krb4_convert = false
}
```

2. Create a KDC database by using the `kdb5_util` utility.

```
# /usr/sbin/kdb5_util -r EXAMPLE.LOCAL create -s
...
```

3. Edit the administrative access control list file (`kadm5.acl`) so that it matches the file below.

```
# cat /etc/krb5/kadm5.acl
#
# Copyright (c) 1998-2000 by Sun Microsystems, Inc.
# All rights reserved.
#
#pragma ident    "@(#)kadm5.acl  1.1      01/03/19 SMI"

*/admin@EXAMPLE.LOCAL *
```

4. Create an administrative principal.

```
# /usr/sbin/kadmin.local
kadmin.local: addprinc root/admin
...
```

5. Create a keytab file (`kadm5.keytab`) for the `kadmind` service.
After the keytab file has been created, exit the `kadmin.local` command.

```
kadmin.local: ktadd -k /etc/krb5/kadm5.keytab kadmin/kdc1.example.local
...
kadmin.local: ktadd -k /etc/krb5/kadm5.keytab changepw/kdc1.example.local
...
kadmin.local: ktadd -k /etc/krb5/kadm5.keytab kadmin/changepw
...
kadmin.local: quit
```

6. Start the Kerberos server daemon.

```
# svcadm enable -r network/security/krb5kdc
# svcadm enable -r network/security/kadmin
```

Note:

If DNS has not been enabled, you cannot start the Kerberos server daemon by executing the `svcadm` command.

7. Acquire the initial ticket for the administrative principal.
Make sure that the initial ticket has been properly acquired.

```
# kinit root/admin
Password for root/admin@EXAMPLE.LOCAL:
# klist -5
Ticket cache: FILE:/tmp/krb5cc_0
Default principal: root/admin@EXAMPLE.LOCAL
```

8. Create a host principal for the KDC server by using the `kadmin` utility.

```
# /usr/sbin/kadmin -p root/admin
...
kadmin: addprinc -randkey host/kdc1.example.local
...
```

9. Add the host principal to the `kadmind` service keytab file (`kadm5.keytab`) by using the `kadmin` utility.

```
kadmin: ktadd host/kdc1.example.local
...
```

10. Create an NFS service principal for each host by using the `kadmin` utility, and then add the NFS service principal to the keytab file.

```
kadmin: addprinc -randkey nfs/node0.example.local
...
kadmin: ktadd -e des-cbc-crc:normal -k /tmp/node0.keytab nfs/
node0.example.local
...
kadmin: addprinc -randkey nfs/node1.example.local
...
kadmin: ktadd -e des-cbc-crc:normal -k /tmp/node1.keytab nfs/
```

```
node1.example.local
...
kadmin: addprinc -randkey nfs/cl1.example.local
...
kadmin: ktadd -e des-cbc-crc:normal -k /tmp/cl1.keytab nfs/
cl1.example.local
...
kadmin: quit
```

For HP-UX 11i v3

For this subsection, the following version of HP-UX 11i v3 is used:

- HP-UX B.11.31 U 9000/800 1801453303 unlimited-user license
- HP-UX11i-OE B.11.31 HP-UX Foundation Operating Environment
- HP-UX11i-OE.OE B.11.31 HP-UX OE control script product

To configure a KDC server on an HP-UX 11i v3 computer and add NFS service principals:

1. Use the `krbsetup` command to create Kerberos settings files (`krb.conf`, `krb.realms`).

The `krbsetup` command enables you to specify values in interactive mode.

```
# /opt/krb5/sbin/krbsetup

Kerberos Server Configuration - Main Menu
-----

Select one of the following options:

1) Configure the Server
2) Start the Kerberos daemons
3) Stop the Kerberos daemons
4) Unconfigure the Server
5) Exit

6) Help

Selection: 1
```

Press the **Enter** key.

```
1) Configure the Server with LDAP backend
2) Configure the Server with C-Tree backend
0) Return to Previous Menu

Selection: [0] 2
```

Press the **Enter** key.

```
1) Configure as a Primary Security Server
2) Configure as a Secondary Security Server
```

```
Selection: 1
```

Press the **Enter** key.

```
1) Configure as a Primary Security Server
2) Configure as a Secondary Security Server
```

```
Selection: 1
```

Press the **Enter** key.

```
THIS MACHINE WILL BE CONFIGURED AS A PRIMARY SERVER

What type of the security mechanism you want to use (DES-MD5/DES-
CRC/DES3)
If you do not select any security mechanism, the default,
DES-MD5 will be selected: DES-CRC
You have selected DES-CRC

Do you want to stash the principal database key
on your local disk (y/n)? [y] :y

Enter the fully qualified name of the Secondary Security Server 1
press 'q' if you want to skip this and proceed further: q

Enter the realm name (the allowed chars are "a-z""A-Z""0-9" "."
"- " _ " *")
If nothing is typed the default name [ KDC1.EXAMPLE.LOCAL ] will be
considered: EXAMPLE.LOCAL

/opt/krb5/krb.conf moved to /opt/krb5/krb.conf.keep
/opt/krb5/krb.realms moved to /opt/krb5/krb.realms.keep
/opt/krb5/kpropd.ini moved to /opt/krb5/kpropd.ini.keep
/etc/krb5.conf moved to /etc/krb5.conf.keep

Creating krb.conf and krb.realms files
Copying admin_acl_file and password.policy file onto /opt/krb5 dir

Do you want to store the log messages in a different directory
rather than
the syslog file (y/n)? [n] : n
You will be prompted for the database Master Password.
It is important that you DO NOT FORGET this password.

Enter Password:
Re-enter Password:

Kerberos server has been configured successfully.

Kerberos daemons are successfully started
Press Enter to go back to the main menu.
```

Press the **Enter** key.

```
Kerberos Server Configuration - Main Menu
-----

Select one of the following options:
```

```
1) Configure the Server
2) Start the Kerberos daemons
3) Stop the Kerberos daemons
4) Unconfigure the Server
5) Exit

6) Help

Selection: 5

You have selected 5 Exiting...
```

2. Check the Kerberos settings files (`krb.conf`, `krb.realms`).

```
# cat /opt/krb5/krb.conf
EXAMPLE.LOCAL
EXAMPLE.LOCAL kdl.example.local admin server
# cat /opt/krb5/krb.realms
*.example.local EXAMPLE.LOCAL
```

3. Edit the administrative access control list file (`admin_acl_file`).
If there is no administrative access control list file, create one.

```
# cat /opt/krb5/admin_acl_file
K/M          CI # needed for kadmd on secondaries
*/admin      * # created by krbsetup can be modified by administrator
```

4. Use the `kadminl` command to create a KDC server `host` principal.

```
# /opt/krb5/admin/kadminl -R "ext host/kdl.example.local"
Connecting as: K/M
Connected to krb5v01 in realm EXAMPLE.LOCAL.
Principal added.
Key extracted.
Disconnected.
```

5. Edit the Kerberos daemon start-up file (`krbsrv`) so that it matches the file below.

```
# cat /etc/rc.config.d/krbsrv
KDC=1
ADMD=1
```

6. Start the Kerberos daemon.

```
# /sbin/init.d/krbsrv start
Starting Kerberos Server Daemons
/opt/krb5/sbin/kdcd
/opt/krb5/sbin/kadmind
Finished startup.

NOTE : If the machine is a primary server please start the kpropd manually.
For more information on propogation refer 'Installing , Configuring HP's
Kerberos server document'
# /opt/krb5/sbin/kpropd
```

7. Use the `kadminl` command to create an NFS service principal for each host, and then add the NFS service principal to the keytab file.

```

# /opt/krb5/admin/kadminl
Connecting as: K/M
Connected to krb5v01 in realm EXAMPLE.LOCAL.
Command: ext
Name of Principal (host/kdc1.example.local): nfs/node0.example.local
Service Key Table File Name (/opt/krb5/v5srvtab): /tmp/node0.keytab
Enter policy name (Press enter key to apply default policy) :
Principal added.
Key extracted.

Command: ext
Name of Principal (host/kdc1.example.local): nfs/node1.example.local
Service Key Table File Name (/opt/krb5/v5srvtab): /tmp/node1.keytab
Enter policy name (Press enter key to apply default policy) :
Principal added.
Key extracted.

Command: ext
Name of Principal (host/kdc1.example.local): nfs/cl1.example.local
Service Key Table File Name (/opt/krb5/v5srvtab): /tmp/cl1.keytab
Enter policy name (Press enter key to apply default policy) :
Principal added.
Key extracted.

Command: q
Disconnected.

```

For AIX 5L V5.3

For this subsection, the following version of AIX 5L V5.3 is used:

- AIX 3 5 000B9B6F4C00
- 5300-09

For Kerberos authentication on an AIX 5L V5.3 computer, install the following file sets ahead of time.

- krb5.client.rte
- modcrypt.base
- clic.rte

To configure a KDC server on an AIX 5L V5.3 computer and add NFS service principals:

1. Use the `mkkrb5srv` command to create a KDC database.

```

# mkkrb5srv -r EXAMPLE.LOCAL -d example.local -s kdc1.example.local
...

```

2. Edit the Kerberos configuration file (`krb5.conf`) so that it matches the file below.

```

# cat /etc/krb5/krb5.conf
[libdefaults]
    default_realm = EXAMPLE.LOCAL
    default_keytab_name = FILE:/etc/krb5/krb5.keytab
    default_tkt_enctypes = des3-cbc-sha1 arcfour-hmac aes256-cts des-
cbc-md5 des-cbc-crc

```

```

    default_tgs_etypes = des3-cbc-sha1 arcfour-hmac aes256-cts des-
    cbc-md5 des-cbc-crc

[realms]
    EXAMPLE.LOCAL = {
        kdc = kdc1.example.local:88
        admin_server = kdc1.example.local:749
        default_domain = example.local
    }

[domain_realm]
    .example.local = EXAMPLE.LOCAL
    kdc1.example.local = EXAMPLE.LOCAL

[logging]
    kdc = FILE:/var/krb5/log/krb5kdc.log
    admin_server = FILE:/var/krb5/log/kadmin.log
    default = FILE:/var/krb5/log/krb5lib.log

```

3. Check the KDC type.

```

# cat /etc/krb5/krb5_cfg_type
master

```

4. Acquire the initial ticket for the administrative principal.

```

# kinit admin/admin@EXAMPLE.LOCAL

```

5. Use the `kadmin.local` command to create a root user principal.

```

# /usr/krb5/sbin/kadmin.local
kadmin.local: addprinc -e des-cbc-crc:normal root
...

```

6. Use the `kadmin.local` command to create a host principal for the KDC user and to add the host principal to the KDC server keytab file (`krb5.keytab`).

After the host principal has been added, exit the `kadmin.local` command.

```

kadmin.local: addprinc -randkey host/kdc1.example.local
...
kadmin.local: ktadd host/kdc1.example.local
...
kadmin.local: q

```

7. Use the `kadmin` utility to create an NFS service principal for each host, and then create a keytab file.

After the keytab file has been created, exit the `kadmin` utility.

```

# /usr/krb5/sbin/kadmin
...
kadmin: add_principal -e des-cbc-crc:normal -randkey nfs/
node0.example.local
...
kadmin: ktadd -e des-cbc-crc:normal -k /tmp/node0.keytab nfs/
node0.example.local
...
kadmin: add_principal -e des-cbc-crc:normal -randkey nfs/
node1.example.local

```

```

...
kadmin: ktadd -e des-cbc-crc:normal -k /tmp/node1.keytab nfs/
node1.example.local
...
kadmin: add_principal -e des-cbc-crc:normal -randkey nfs/cl1.example.local
...
kadmin: ktadd -e des-cbc-crc:normal -k /tmp/cl1.keytab nfs/
cl1.example.local
...
kadmin: q

```

8. Restart the KDC server.

```

# /usr/krb5/sbin/stop.krb5
...
# /usr/krb5/sbin/start.krb5
...

```

Distributing and retrieving keytab files

This section explains how to distribute a keytab file in which all of the NFS service principals for each host are added by merging the keytab files managed by each host. For details on creating a keytab file that will be distributed to each host, see [Configuring the KDC server and adding NFS service principals on page C-3](#).

Keytab file distribution destinations

For this example, the following table lists the keytab file distribution destinations that are used when configuring the KDC server.

Table C-3 Keytab file distribution destinations

#	Keytab file name	Target host	Distribution destination
1	node0.keytab	Node 0	Node 0: /home/nasroot Node 1: /home/nasroot#
2	node1.keytab	Node 1	Node 0: /home/nasroot# Node 1: /home/nasroot
3	cl1.keytab	NFS client	NFS client: /tmp

#

Distribute the keytab file to both nodes so that you can continue operations even if a failover occurs.

Distributing keytab files

A keytab file includes confidential information. While always keeping data security, please distribute keytab files in the following way.

For Windows:

Transfer keytab files by using software that enables the keytab files to be safely replicated.

For UNIX:

Transfer keytab files by using `scp`.

Retrieving keytab files (for HDI nodes)

To retrieve distributed keytab files on HDI nodes:

1. Use the `nfskeytabadd` command to merge the keytab files.
Execute the `nfskeytabadd` command on both nodes.

```

$ sudo nfskeytabadd -i /home/nasroot/node0.keytab
$ sudo nfskeytabadd -i /home/nasroot/node1.keytab
$ sudo nfskeytablist
slot KVNO Principal
-----
1      3      nfs/node0.example.local@EXAMPLE.LOCAL
2      3      nfs/node1.example.local@EXAMPLE.LOCAL

```

Retrieving keytab files (for an NFS client)

For this example, one of the following platforms is being used on the NFS client.

Table C-4 Platforms used on the NFS client

#	Platform	Version
1	Red Hat Enterprise Linux Advanced Platform v5.6	Linux version 2.6.18-238.el5
		Red Hat Enterprise Linux Server release 5 (Tikanga)
2	Solaris 10	SunOS 5.10 Generic_137137-09 sun4u sparc SUNW, Sun-Blade-1000
		Solaris 10 10/08 s10s_u6wos_07b SPARC Copyright 2008 Sun Microsystems, Inc. All Rights Reserved. Use is subject to license terms. Assembled 27 October 2008
3	HP-UX 11i v3	HP-UX B.11.31 U 9000/800 1801453303 unlimited-user license
		HPUX11i-OE B.11.31 HP-UX Foundation Operating Environment HPUX11i-OE.OE B.11.31 HP-UX OE control script product
4	AIX 5L V5.3	AIX 3 5 000B9B6F4C00
		5300-09

The procedure for retrieving keytab files on NFS clients is different for AIX and all other platforms.

To retrieve distributed keytab files on NFS clients that are using a platform other than AIX:

1. Edit the Kerberos configuration file (`krb5.conf`).
Change the KDC server domain name and KDC server name.

```
# cat /etc/krb5.conf
[logging]
default = FILE:/var/log/krb5libs.log
kdc = FILE:/var/log/krb5kdc.log
admin_server = FILE:/var/log/kadmind.log

[libdefaults]
default_realm = EXAMPLE.LOCAL
dns_lookup_realm = false
dns_lookup_kdc = false
ticket_lifetime = 24h
forwardable = yes

[realms]
EXAMPLE.LOCAL = {
    kdc = kdcl.example.local:88
    admin_server = kdcl.example.local:749
    default_domain = example.local
}

[domain_realm]
.example.local = EXAMPLE.LOCAL
example.local = EXAMPLE.LOCAL

[kdc]
profile = /var/kerberos/krb5kdc/kdc.conf

[appdefaults]
pam = {
    debug = false
    ticket_lifetime = 36000
    renew_lifetime = 36000
    forwardable = true
    krb4_convert = false
}
```

2. Use the `ktutil` command to merge the keytab files.
Specify the keytab file managed by an NFS client. For this example, `/etc/krb5.keytab` is specified.

```
# ktutil
ktutil: rkt /tmp/cl1.keytab
ktutil: wkt /etc/krb5.keytab
ktutil: list
slot KVNO Principal
----
-----
1 3 nfs/cl1.example.local@EXAMPLE.LOCAL
ktutil: quit
```

To retrieve distributed keytab files on NFS clients that are using AIX 5L V5.3:

1. Make sure that the `krb5.client.rte`, `modcrypt.base`, and `clie.rte` file sets has been installed.

```

# lslpp -l krb5.client.rte
Fileset                                Level State      Description
-----
-
Path: /usr/lib/objrepos
  krb5.client.rte                      1.4.0.8 COMMITTED  Network Authentication
Service Client

Path: /etc/objrepos
  krb5.client.rte                      1.4.0.8 COMMITTED  Network Authentication
Service Client
# lslpp -l | grep modcrypt.base
modcrypt.base.includes                5.3.7.1 COMMITTED  Cryptographic Library
Include
modcrypt.base.lib                     5.3.7.1 COMMITTED  Cryptographic Library
# lslpp -l | grep clic.rte
clic.rte.includes                    3.24.0.1 COMMITTED  CryptoLite for C Library
clic.rte.kernext                     3.24.0.1 COMMITTED  CryptoLite for C Kernel
clic.rte.lib                          3.24.0.1 COMMITTED  CryptoLite for C Library
clic.rte.kernext                     3.24.0.1 COMMITTED  CryptoLite for C Kernel

```

2. Use the `config.krb5` command to configure the Kerberos client.

```

# config.krb5 -C -d example.local -r EXAMPLE.LOCAL -c kdc1.example.local -
s kdc1.example.local

```

3. Edit the Kerberos configuration file (`krb5.conf`).

Change the KDC server domain name and KDC server name.

```

# cat /etc/krb5/krb5.conf
[libdefaults]
    default_realm = EXAMPLE.LOCAL
    default_keytab_name = FILE:/etc/krb5/krb5.keytab
    default_tkt_enctypes = des3-cbc-sha1 arcfour-hmac aes256-cts des-
cbc-md5 des-cbc-crc
    default_tgs_enctypes = des3-cbc-sha1 arcfour-hmac aes256-cts des-
cbc-md5 des-cbc-crc

[realms]
    EXAMPLE.LOCAL = {
        kdc = kdc1.example.local:88
        admin_server = kdc1.example.local:749
        default_domain = example.local
    }

[domain_realm]
    .example.local = EXAMPLE.LOCAL
    kdc1.example.local = EXAMPLE.LOCAL

[logging]
    kdc = FILE:/var/krb5/log/krb5kdc.log
    admin_server = FILE:/var/krb5/log/kadmin.log
    default = FILE:/var/krb5/log/krb5lib.log

```

4. Use the `ktutil` command to merge the keytab files.

```

# /usr/krb5/sbin/ktutil
ktutil: rkt /tmp/cl1.keytab
ktutil: wkt /etc/krb5/krb5.keytab
ktutil: list
slot KVNO Principal
----
-----

```

```
1 3      nfs/cl1.example.local@EXAMPLE.LOCAL
ktutil: quit
```

5. Set up the `gssd` daemon so that the merged keytab file can be used.

```
# nfshostkey -p nfs/aix.example.local -f /etc/krb5/krb5.conf
```

6. Start the `gssd` daemon.

```
# chnfs -S -B
```



Accessing NFS shared directories when Kerberos authentication is used

This section explains (through the use of execution examples) how to access NFS shared directories when Kerberos authentication is used.

- [Specifying a security flavor from File Services Manager](#)
- [Mounting shared directories from NFS clients](#)
- [Accessing NFS shared directories](#)

Specifying a security flavor from File Services Manager

You can specify which security flavors to use for Kerberos authentication whenever you create NFS shares or change the NFS service configuration from File Services Manager.

The following figure shows an example of specifying security flavors when an NFS share is created from the File Services Manager GUI.

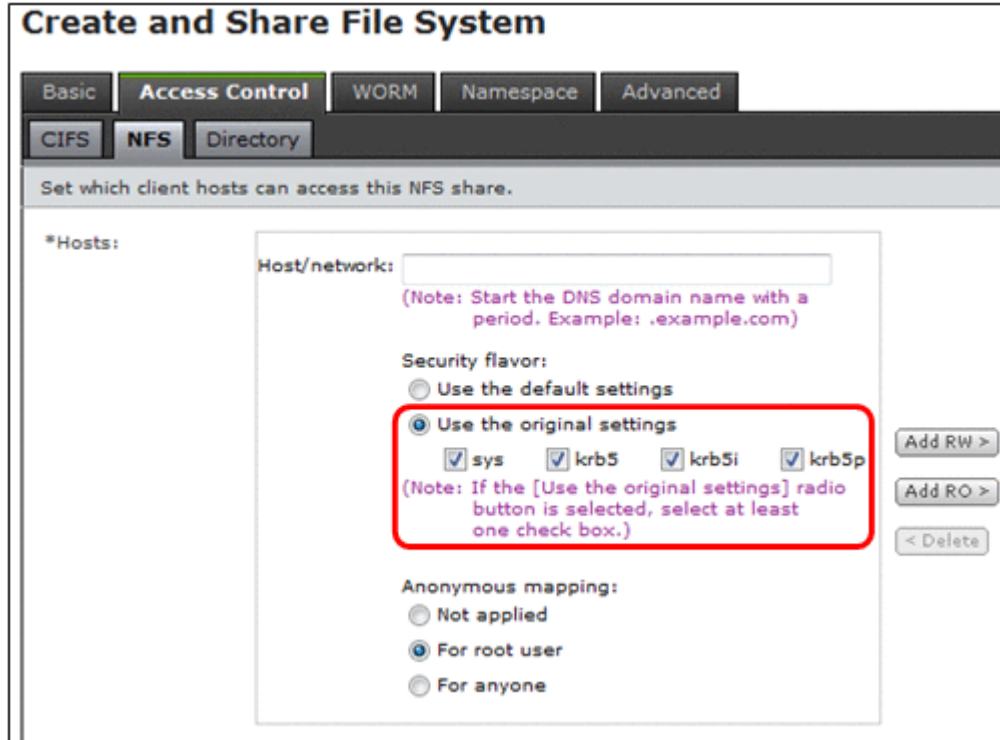


Figure D-1 Example of specifying security flavors

For details on how to specify security flavors when creating NFS shares or changing the NFS service configuration, see the *Administrator's Guide*.

Mounting shared directories from NFS clients

When mounting a shared directory from an NFS client, specify one security flavor from among those that have been set up in File Services Manager.

The following execution examples show (for all of the security flavors) how to mount the shared directory `node0.example.local:/mnt/nfs01` from an NFS client with Red Hat by using the NFSv3 protocol.

- When using Kerberos 5:

```
# mount -o vers=3,sec=krb5 node0.example.local:/mnt/nfs01 /mnt
# mount
:
node0.example.local:/mnt/nfs01 on /mnt type nfs
(rw,sec=krb5,addr=192.168.0.10)
```

- When using Kerberos 5 (Integrity):

```
# mount -o vers=3,sec=krb5i node0.example.local:/mnt/nfs01 /mnt
# mount
:
node0.example.local:/mnt/nfs01 on /mnt type nfs
(rw,sec=krb5i,addr=192.168.0.10)
```

- When using Kerberos 5 (Privacy):

```
# mount -o vers=3,sec=krb5p node0.example.local:/mnt/nfs01 /mnt
# mount
:
node0.example.local:/mnt/nfs01 on /mnt type nfs
(rw,sec=krb5p,addr=192.168.0.10)
```

- When using AUTH_SYS:

```
# mount -o vers=3,sec=sys node0.example.local:/mnt/nfs01 /mnt
# mount
:
node0.example.local:/mnt/nfs01 on /mnt type nfs
(rw,sec=sys,addr=192.168.0.10)
```

Accessing NFS shared directories

To access mounted NFS shared directories as root and general users, you need to assign both the root and user principals to the KDC server domain. If the KDC server has been configured on a Windows server, a root or general user who wants to access the NFS shared directories needs to be registered as an Active Directory user.

Users who have acquired initial tickets can access the NFS shared directories.

A ticket is generally valid for 8-10 hours. Be sure to also consider the fact that the file systems might be used for time-consuming batch processes, in which case, the KDC policy settings will need to be changed.



Adding a secondary KDC server

You can have up to five KDC servers in an HDI system. When adding another (secondary) KDC server, you need to replicate the KDC database on the primary KDC server to the secondary KDC server.

- [Procedure for adding a KDC server](#)

Procedure for adding a KDC server

This section explains how to configure and add the secondary KDC server `kdc2.example.local` on a Red Hat Enterprise Linux Advanced Platform v5.2 computer. For details on the prerequisites for configuring KDC servers, see [Before configuring the KDC server on page C-3](#).

1. Make sure that the `krb5-server`, `krb5-libs`, and `krb5-workstation` packages have been installed.

```
# rpm -qa | grep krb
krb5-server-1.5-17
krb5-libs-1.5-17
krb5-workstation-1.5-17
```

2. Use the `kdb5_util` utility to create a KDC database on the secondary KDC server.

The KDC database must be created in the same way it was created on the primary KDC server.

```
# /usr/kerberos/sbin/kdb5_util -r EXAMPLE.LOCAL create -s
Loading random data
Initializing database '/var/kerberos/krb5kdc/principal' for realm
'EXAMPLE.LOCAL',
master key name 'K/M@EXAMPLE.LOCAL'
You will be prompted for the database Master Password.
It is important that you NOT FORGET this password.
Enter KDC database master key:
Re-enter KDC database master key to verify:
```

3. Edit the Kerberos configuration file (`krb5.conf`).

The Kerberos configuration file must be the same on the primary the secondary KDC servers.

```
# cat /etc/krb5.conf
[logging]
default = FILE:/var/log/krb5libs.log
kdc = FILE:/var/log/krb5kdc.log
admin_server = FILE:/var/log/kadmind.log

[libdefaults]
default_realm = EXAMPLE.LOCAL
dns_lookup_realm = false
dns_lookup_kdc = false
ticket_lifetime = 24h
forwardable = yes

[realms]
EXAMPLE.LOCAL = {
kdc = kdc1.example.local:88
kdc = kdc2.example.local:88
admin_server = kdc1.example.local:749
default_domain = example.local
}

[domain_realm]
.example.local = EXAMPLE.LOCAL
example.local = EXAMPLE.LOCAL

[kdc]
profile = /var/kerberos/krb5kdc/kdc.conf
```

```
[appdefaults]
pam = {
    debug = false
    ticket_lifetime = 36000
    renew_lifetime = 36000
    forwardable = true
    krb4_convert = false
}
```

4. Use the `kadmin` utility to create a host principal on the secondary KDC server.

```
# kadmin
Password for root/admin@EXAMPLE.LOCAL:
kadmin: add_principal -randkey host/kdc2.example.local
WARNING: no policy specified for host/kdc2.example.local@EXAMPLE.LOCAL;
defaulting to no policy
Principal "host/kdc2.example.local@EXAMPLE.LOCAL" created.
```

5. Use the `kadmin` utility to add the host principal to the keytab file (`krb5.keytab`) on the secondary KDC server.

Make sure that the host principal has been correctly added.

```
kadmin: ktadd host/kdc2.example.local
Entry for principal host/kdc2.example.local with kvno 3, encryption type
Triple DES cbc mode with HMAC/sha1 added to keytab WRFILE:/etc/krb5.keytab.
Entry for principal host/kdc2.example.local with kvno 3, encryption type
DES cbc mode with CRC-32 added to keytab WRFILE:/etc/krb5.keytab.
```

6. Create and edit the `kpropd.acl` file.
Create the `kpropd.acl` file in the directory `/var/kerberos/krb5kdc`, which is where the KDC database is stored. Also, add the host principals for all of the secondary KDC servers that are part of the KDC server domain.

```
# cat /var/kerberos/krb5kdc/kpropd.acl
host/kdc1.example.local@EXAMPLE.LOCAL
host/kdc2.example.local@EXAMPLE.LOCAL
```

7. Start the `kpropd` daemon on the primary and secondary KDC servers.

```
# kpropd -S
```

8. Transfer a copy of the KDC database that was dumped from the primary KDC server to the secondary KDC server.

Using `cron` allows you to execute this step on a regular schedule.

```
# /usr/kerberos/sbin/kdb5_util dump /var/kerberos/krb5kdc/slave_datatrans
# kprop -d -f /var/kerberos/krb5kdc/slave_datatrans kdc2.example.local
3310 bytes sent.
Database propagation to kdc2.example.local: SUCCEEDED
```

9. Create a stash file on the secondary KDC server.
The KDC database master key is stored in the stash file.

```
# /usr/kerberos/sbin/kdb5_util stash
Enter KDC database master key:
```

10. Start the Kerberos server daemon on the secondary KDC server.

```
# /usr/kerberos/sbin/krb5kdc
```



APIs for WORM operation

To create a WORM file in a WORM file system, you must set a retention period (storage period) for the file to make it read-only. To set or extend a retention period for a file, use a custom application that you have created. This section introduces the APIs for creating such custom applications.

- [Creating a WORM file from a CIFS share file](#)
- [Creating a WORM file from an NFS share file](#)

Creating a WORM file from a CIFS share file

To create a WORM file from a CIFS share file, use a Windows API.

Creating a WORM file

The procedure for creating a WORM file is described below.

To create a WORM file:

1. Create a writable file and write data to it.
2. Set a retention period for the file.
3. Make the file read-only.

APIs required for creating WORM files

The following table lists the Windows APIs required for creating WORM files.

Table F-1 APIs required for creating WORM files (CIFS share)

Function name	Description	Reference documents
SetFileTime	Sets a retention period for a file.	http://msdn.microsoft.com/en-us/library/ms724933(VS.85).aspx
SetFileAttributes	Makes a file read-only or clears the read-only attribute.	http://msdn.microsoft.com/en-us/library/aa365535(VS.85).aspx

SetFileTime

This section explains SetFileTime.

Name

```
SetFileTime
```

Format

```
BOOL SetFileTime(  
    HANDLE          hFile,           // File handle  
    CONST FILETIME *lpCreationTime,  // Creation date to be set  
    CONST FILETIME *lpLastAccessTime, // Access date to be set  
    CONST FILETIME *lpLastWriteTime  // Update date to be set  
);
```

Functional description

This API updates the timestamp of the specified file.

Arguments

Neither `lpCreationTime` nor `lpLastWriteTime` is required for creating WORM files. Specify `NULL` for these parameters (which indicates that the corresponding time-stamp will not be updated).

The `FILETIME` type is not suitable for direct interactive handling by the user. We therefore recommend using a program that converts data of the

SYSTEMTIME type to the FILETIME type. The following table describes the FILETIME type and the SYSTEMTIME type structures.

Table F-2 FILETIME type and SYSTEMTIME type structures

Structure name	Members	Description	Reference documents
FILETIME	DWORD dwLowDateTime; DWORD dwHighDateTime;	A 64-bit value indicating the number of 100-nanosecond intervals since January 1, 1601. This type is required as an argument of <code>SetFileTime</code> , but it is not suitable for direct interactive handling by a user.	http://msdn.microsoft.com/en-us/library/ms724284(VS.85).aspx
SYSTEMTIME	WORD wYear; WORD wMonth; WORD wDay; WORD wDayOfWeek; WORD wHour; WORD wMinute; WORD wSecond; WORD wMilliseconds;	Uses each member to indicate the time consisting of the year, month, date, day of the week, hour, minute, second, and millisecond.	http://msdn.microsoft.com/en-us/library/ms724950(VS.85).aspx

SetFileAttributes

This section explains `SetFileAttributes`.

Name

`SetFileAttributes`

Format

```
BOOL SetFileAttributes (
    LPCTSTR lpFileName, // File name
    DWORD dwFileAttributes // Attribute to be set
);
```

Functional description

This API sets the DOS attribute of the specified file.

Arguments

Suppose you want to add a specific attribute to the existing attributes of a file. To do this, you must obtain the existing attributes from the target

file, and then specify the obtained attributes and the additional attribute value in `dwFileAttributes`.

Useful APIs for creating WORM files

The following table lists a number of useful Windows APIs that can be used with a program for creating WORM files.

Table F-3 Useful APIs for creating WORM files

Function name	Description	Reference documents
<code>SystemTimeToFileTime</code>	Converts data of the <code>SYSTEMTIME</code> type to the <code>FILETIME</code> type, which is supported by <code>SetFileTime</code> .	http://msdn.microsoft.com/en-us/library/ms724948(VS.85).aspx
<code>LocalFileTimeToFileTime</code>	Converts the local time to Coordinated Universal Time (UTC).	http://msdn.microsoft.com/en-us/library/ms724490(VS.85).aspx
<code>CreateFile</code>	Obtains the handle of the file to be specified in <code>SetFileTime</code> .	http://msdn.microsoft.com/en-us/library/aa363858(VS.85).aspx
<code>GetFileAttributes</code>	Obtains the existing file attributes.	http://msdn.microsoft.com/en-us/library/aa364944(VS.85).aspx

Sample program

Below is a sample C program that sets a retention period for a file, and then makes the file read-only.

```
#include <windows.h>
#include <stdio.h>
#include <string.h>

void getTimestamp(FILETIME *ftLpTime, char *tcArgtime)
{
    SYSTEMTIME stFileTime;
    FILETIME ftLocalFileTime;

    /*Convert input value to SYSTEMTIME type*/
    memset(&stFileTime, 0, sizeof(SYSTEMTIME));
    sscanf(tcArgtime, "%d/%d/%d %d:%d:%d",
           &(stFileTime.wYear), &(stFileTime.wMonth),
           &(stFileTime.wDay), &(stFileTime.wHour),
           &(stFileTime.wMinute), &(stFileTime.wSecond)
    );
    stFileTime.wMilliseconds = 0;

    /*Convert from SYSTEMTIME type to FILETIME type */
    SystemTimeToFileTime(&stFileTime, &ftLocalFileTime);
    /*Convert local time to Coordinated Universal Time (UTC)*/
    LocalFileTimeToFileTime(&ftLocalFileTime, ftLpTime);
}

int main(int argc, char *argv[])
{
    char *filename;
```



```

char *filetime;
HANDLE h;
FILETIME ftLastAccessTime;
DWORD attr;

/*Check arguments*/
if (argc != 3) {
    fprintf(stderr, "usage: %s time file \n", argv[0]);
    fprintf(stderr, "          ex. time: \"2040/12/31 23:59:59\\n\\n");
    return 1;
}
filetime = argv[1];
filename = argv[2];

/*Obtain the file handle*/
h = CreateFile(
    filename, FILE_WRITE_ATTRIBUTES, 0, NULL,
    OPEN_EXISTING, FILE_FLAG_BACKUP_SEMANTICS, NULL
);
if (h == INVALID_HANDLE_VALUE) {
    fprintf(stderr, "CreateFile error: ");
    return 1;
}

/*Set retention period for the file*/
getTimestamp(&ftLastAccessTime, filetime);
if (!SetFileTime(h, NULL, &ftLastAccessTime, NULL)) {
    fprintf(stderr, "SetFileTime error: ");
    CloseHandle(h);
    return 1;
}
CloseHandle(h);

/*Grant the read-only attribute to the file*/
attr = GetFileAttributes(filename);
attr |= FILE_ATTRIBUTE_READONLY;
if (!SetFileAttributes(filename, attr)) {
    fprintf(stderr, "SetFileAttributes error: ");
    return 1;
}

return 0;
}

```

The following is an execution example of a sample program. Below that is an example of the file's properties before and after the execution of the sample program. For these examples, 12:00:00 on January 30, 2015 is set for the retention period (storage period) for the file.

```

\\10.213.88.155\worm_ad\sample_dir\worm.exe "2015/1/30 12:00:00" \
10.213.88.155\worm_ad\sample_dir\worm_file.doc

```

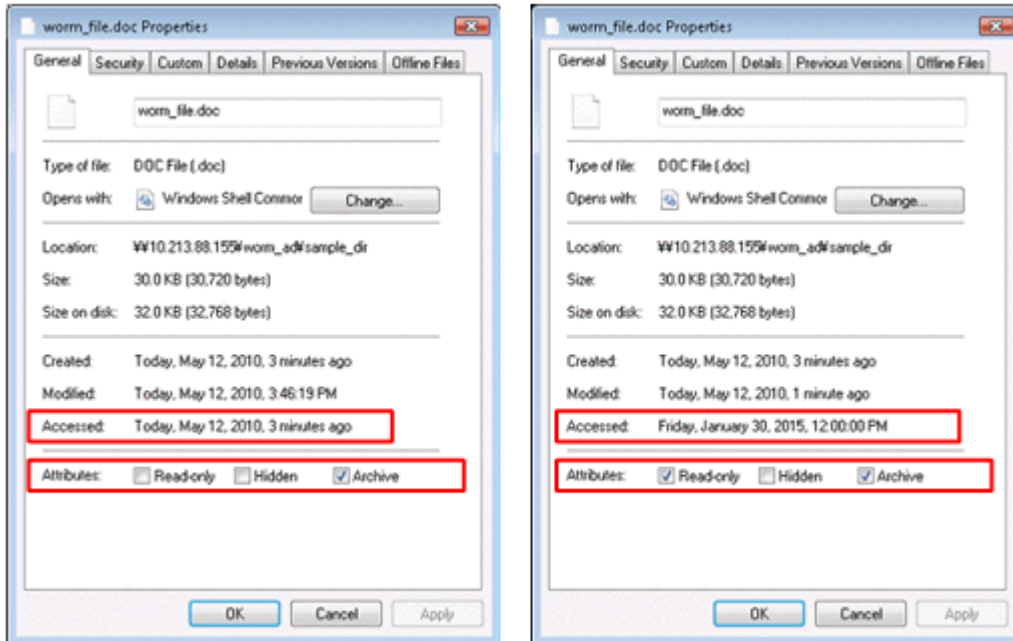


Figure F-1 Example of file properties before and after execution of the sample program (left: before execution, right: after execution)

Creating a WORM file from an NFS share file

To create a WORM file from an NFS share file, use system calls.

Creating a WORM file

The procedure for creating a WORM file is described below.

To create a WORM file:

1. Create a writable file and write data to it.
2. Set a retention period for the file.
3. Make the file read-only.

APIs required for creating WORM files

The following table lists the API system calls required for creating WORM files.

Table F-4 APIs required for creating WORM files (NFS share)

System call	Description
utime() utimes()	Sets a retention period for a file.
chmod() fchmod()	Makes a file read-only, or clears the read-only attribute.

utime(), utimes()

This section explains `utime()` and `utimes()`.

Name

`utime`
`utimes`

Format

```
#include <sys/types.h>
#include <utime.h>
int utime(const char *filename, const struct utimbuf *times);

#include <sys/time.h>
int utimes(const char *filename, const struct timeval times[2]);
```

Functional description

This system call changes the last access time (`atime`) and modified time (`mtime`) of a specified file.

Arguments

Set the retention period for the value of `atime`, and set the file's current setting for the value of `mtime`. Changing both `atime` and `mtime` at once is interpreted not as a retention period change but as a file attribute change. In this case, a system call error occurs. A system call that changes the value of `atime` for a WORM file might be processed as a system call for changing the retention period.

The following example shows how to define the `utimbuf` structure:

```
struct utimbuf {
    time_t actime;    //Set the retention period
    time_t modtime;  //Set the current value in the file
};
```

chmod(), fchmod()

This section explains `chmod()` and `fchmod()`.

Name

`chmod`
`fchmod`

Format

```
#include <sys/stat.h>
int chmod(const char *path, mode_t mode);

int fchmod(int fd, mode_t mode);
```

Functional description

These system calls change the permission for the specified file.

Arguments

To make the file read-only, set the write permission for `S_IWUSR` (owner), `S_IWGRP` (group to which the owner belongs), and `S_IWOTH` (other users) to off. To clear the read-only attribute, set the write permission for

`S_IWUSR`, `S_IWGRP`, or `S_IWOTH` to on. The settings for the read permission and execute permission cannot be modified.

Sample program

Below is a sample program that sets a retention period for a file, and then makes the file read-only.

Below is an example of a program that specifies the target file in the first argument and the retention period in the second argument, and then changes the file to a WORM file.

Specify the retention period using the current time as the reference value. For example, to specify a retention period of 300 seconds from the current time, specify `300`. You can specify whether the retention period is in units of days, months, or years by specifying `d`, `m`, or `y` after the numerical value. You must coordinate the times of the NFS client and an HDI node before executing the program.

```
#include <stdio.h>
#include <sys/types.h>
#include <utime.h>
#include <sys/stat.h>
#include <unistd.h>
#include <stdlib.h>

typedef enum { false = 0, true = 1 } boolean;

void
usage (char *cmd)
{
    printf ("usage: %s regular-file retention-time\n", cmd);
    printf ("      retention-time format:\n");
    printf ("      <numbers>d\tdays\n");
    printf ("      <numbers>m\tmonth\n");
    printf ("      <numbers>y\tyear\n");
    printf ("      <numbers>\tsecond\n");
}

time_t
set_worm_file(char *path, time_t retention_time)
{
    struct stat      st;
    struct utimbuf   utim;
    mode_t          new_mod;

    // Obtain the file's current atime and mtime values
    if (stat (path, &st) == -1) {
        return 0;
    }

    // Set the retention period (do not change mtime)
    utim.modtime = st.st_mtime;
    utim.actime  = retention_time;

    if (utime (path, &utim) == -1) {
        return 0;
    }

    // Change the file permission to read-only
    new_mod = (st.st_mode & ~(S_IWUSR | S_IWGRP | S_IWOTH));
    if (chmod (path, new_mod) == -1) {
```

```

        return 0;
    }

    if (stat (path, &st) == -1) {
        return 0;
    }

    return st.st_atime;        // success(return current access time).
}

boolean
is_file (char *path)
{
    struct stat    st;

    if (stat(path, &st) == -1) {
        return false;
    }

    if (S_ISREG(st.st_mode)) {
        return true;
    }

    return false;
}

time_t
convert_time (char *s)
{
    int    value;
    time_t  retval;
    time_t  now_time = time(NULL);

    if (sscanf (s, "%d", &value) == 1) {
        while (*s != '\0') {
            if (!isdigit (*s)) {
                break;
            }
            s++;
        }
        switch (*s) {
            case 'd':
            case 'D':
                printf ("unit is day. (%d)\n", value);
                value = (value * 24 * 3600);
                break;

            case 'm':
            case 'M':
                printf ("unit is month. (%d)\n", value);
                value = (value * 24 * 3600 * 30);
                break;

            case 'y':
            case 'Y':
                printf ("unit is year. (%d)\n", value);
                value = (value * 24 * 3600 * 30 * 365);
                break;

            default:
                printf ("unit is second. (%d)\n", value);
                break;
        }
    }

    retval = (time_t)value + now_time;
}

```

```

int
main (int ac, char **av)
{
    time_t  result;
    time_t  new_atime;

    if (ac < 3) {
        usage (av[0]);
        exit (0);
    }

    if (!is_file (av[1])) {
        usage (av[0]);
        exit (0);
    }

    // setting time information.
    new_atime = convert_time (av[2]);

    // change file to WORM
    result = set_worm_file (av[1], new_atime);

    // Display the retention period (if 0 is displayed, the file was not
    changed to WORM)
    printf ("new access time (%u)\n", result);

    return 0;
}

```

Below is an example of a program that changes a file named `file01` to a WORM file for a retention period of 600 seconds. For this example, it is assumed that `file01` is not 0 bytes.

```

$ ./worm file01 600
unit is second. (600)
now time = 1264843082
new access time (1264843682)

```

WORM errors and system calls when a WORM file is accessed

The following table shows the relationship between WORM-specific errors that might be returned to the NFS client and system calls from the client after a WORM file has been accessed.

Table F-5 Relationship between system calls related to WORM files and errors at access

Protocol version	Procedure/Operation	NFS error	System call from the client
2	NFSPROC_SETATTR	NFSERR_ACCES NFSERR_IO	<code>chmod</code> or <code>utime</code> system call Returns <code>NFSERR_IO (EIO)</code> if an error occurred during a <code>utime</code> system call.
2	NFSPROC_LOOKUP	NFSERR_ACCES	General system calls for referencing a file (such as the <code>open</code> system call)

Protocol version	Procedure/ Operation	NFS error	System call from the client
2	NFSPROC_WRITE	NFSERR_ACCES	write system call
2	NFSPROC_CREATE	NFSERR_ACCES	creat system call
2	NFSPROC_REMOVE	NFSERR_ROFS	unlink system call
2	NFSPROC_RENAME	NFSERR_ACCES NFSERR_IO	rename system call Returns NFSERR_IO for directory rename operations. If the option for renaming empty directories is enabled, empty directories can be renamed.
3	NFS3PROC_SETATTR	NFS3ERR_ACCES NFS3ERR_IO	chmod or utime system call Returns NFS3ERR_IO if an error occurred during a utime system call.
3	NFS3PROC_LOOKUP	NFS3ERR_ACCES	General system calls for referencing a file (such as the open system call)
3	NFS3PROC_WRITE	NFS3ERR_ACCES	write system call
3	NFS3PROC_CREATE	NFS3ERR_ACCES	creat system call
3	NFS3PROC_REMOVE	NFS3ERR_ROFS	unlink system call
3	NFS3PROC_RENAME	NFS3ERR_ACCES NFS3ERR_IO	rename system call Returns NFS3ERR_IO for directory rename operations. If the option for renaming empty directories is enabled, empty directories can be renamed.
3	NFS3PROC_COMMIT	NFS3ERR_IO	write or close system call
4	OP_CLOSE	NFS4ERR_IO	close system call
4	OP_COMMIT	NFS4ERR_IO	write or close system call
4	OP_CREATE	NFS4ERR_ACCESS	creat system call
4	OP_OPEN	NFS4ERR_ACCESS	open system call
4	OP_REMOVE	NFS4ERR_ROFS	unlink system call
4	OP_RENAME	NFS4ERR_ACCESS NFS4ERR_IO	rename system call Returns NFS4ERR_IO for directory rename operations. If the option for renaming empty directories is enabled, empty directories can be renamed.
4	OP_SETATTR	NFS4ERR_ACCESS NFS4ERR_IO	chmod or utime system call

Protocol version	Procedure/ Operation	NFS error	System call from the client
			Returns NFS4ERR_IO if an error occurred during a utime system call.
4	OP_WRITE	NFS4ERR_ACCESS	write system call

The value shown in the NFS error column is returned to the client application. However, the error number differs depending on the version of the protocol that was used for access, so substitute the value given in the following table.

Table F-6 Substitute value for error numbers

Error number	Substitute value
NFSERR_ACCES NFS3ERR_ACCES NFS4ERR_ACCES	EACCES
NFSERR_IO NFS3ERR_IO NFS4ERR_IO	EIO
NFSERR_ROFS NFS3ERR_ROFS NFS4ERR_ROFS	EROFS

Note:

Depending on the client, a different error number might be returned.



References

This appendix contains related Web sites that provide reference information.

- [Web sites](#)

Web sites

Related Web sites and their URLs are as follows:

OpenLDAP

<http://www.openldap.org/>

ADAM-Check the Microsoft website for information on ADAM, or go to:

[http://technet.microsoft.com/en-us/library/cc736765\(W.S.10\).aspx](http://technet.microsoft.com/en-us/library/cc736765(W.S.10).aspx)



Acronyms

This appendix lists the acronyms used in the HDI manuals.

- [Acronyms used in the HDI manuals.](#)

Acronyms used in the HDI manuals.

The following acronyms are used in the HDI manuals.

ABE	Access Based Enumeration
ACE	access control entry
ACL	access control list
AES	Advanced Encryption Standard
AJP	Apache JServ Protocol
API	application programming interface
ARP	Address Resolution Protocol
ASCII	American Standard Code for Information Interchange
ASN	Abstract Syntax Notation
BDC	Backup Domain Controller
BMC	baseboard management controller
CA	certificate authority
CHA	channel adapter
CHAP	Challenge-Handshake Authentication Protocol
CIFS	Common Internet File System
CIM	Common Information Model
CLI	command line interface
CPU	central processing unit
CSR	certificate signing request
CSV	comma-separated values
CTL	controller
CU	control unit
CV	custom volume
DAACL	discretionary access control list
DAR	Direct Access Recovery
DB	database
DBMS	database management system
DC	domain controller
DDNS	Dynamic Domain Name System
DEP	data execution prevention
DES	Data Encryption Standard
DFS	distributed file system
DIMM	dual in-line memory module

DHCP	Dynamic Host Configuration Protocol
DLL	dynamic-link library
DN	distinguished name
DNS	Domain Name System
DOM	Document Object Model
DOS	Disk Operating System
DRAM	dynamic random access memory
DSA	digital signal algorithm
DTD	Document Type Definition
ECC	error-correcting code
EUC	Extended UNIX Code
FC	Fibre Channel
FC-SP	Fibre Channel - Security Protocol
FIB	forwarding information base
FIFO	First In, First Out
FQDN	fully qualified domain name
FTP	File Transfer Protocol
FV	Fixed Volume
FXP	File Exchange Protocol
GbE	Gigabit Ethernet
GID	group identifier
GMT	Greenwich Mean Time
GPL	GNU General Public License
GUI	graphical user interface
HBA	host bus adapter
H-LUN	host logical unit number
HPFS	High Performance File System
HSSO	HiCommand single sign-on
HTML	HyperText Markup Language
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
I/O	input/output
ICAP	Internet Content Adaptation Protocol
ICMP	Internet Control Message Protocol
ID	identifier

IP	Internet Protocol
IP-SW	IP switch
JDK	Java Development Kit
JIS	Japanese Industrial Standards
JSP	JavaServer Pages
KDC	Key Distribution Center
LACP	Link Aggregation Control Protocol
LAN	local area network
LBA	logical block addressing
LCD	Local Configuration Datastore
LDAP	Lightweight Directory Access Protocol
LDEV	logical device
LDIF	LDAP Data Interchange Format
LDKC	logical disk controller
LED	light-emitting diode
LF	Line Feed
LTS	long term support
LU	logical unit
LUN	logical unit number
LUSE	logical unit size expansion
LVI	Logical Volume Image
LVM	Logical Volume Manager
MAC	Media Access Control
MD5	Message-Digest algorithm 5
MIB	management information base
MMC	Microsoft Management Console
MP	microprocessor
MSS	maximum segment size
MTU	maximum transmission unit
NAS	Network-Attached Storage
NAT	network address translation
NDMP	Network Data Management Protocol
NetBIO S	Network Basic Input/Output System
NFS	Network File System

NIC	network interface card
NIS	Network Information Service
NTFS	New Technology File System
NTP	Network Time Protocol
OID	object identifier
ORB	object request broker
OS	operating system
PAP	Password Authentication Protocol
PC	personal computer
PCI	Peripheral Component Interconnect
PDC	Primary Domain Controller
PDU	protocol data unit
PID	process identifier
POSIX	Portable Operating System Interface for UNIX
PP	program product
RADIUS	Remote Authentication Dial In User Service
RAID	Redundant Array of Independent Disks
RAM	random access memory
RAS	Reliability Availability Serviceability
RCS	Revision Control System
RD	relational database
RFC	Request for Comments
RID	relative identifier
RPC	remote procedure call
RSA	Rivest, Shamir, and Adleman
SACL	system access control list
SAN	storage area network
SAS	Serial Attached SCSI
SATA	serial ATA
SAX	Simple API for XML
SCSI	Small Computer System Interface
SFTP	SSH File Transfer Protocol
SHA	secure hash algorithm
SID	security identifier
SJIS	Shift JIS

SLPR	Storage Logical Partition
SMB	Server Message Block
SMD5	Salted Message Digest 5
SMTP	Simple Mail Transfer Protocol
SNMP	Simple Network Management Protocol
SOAP	Simple Object Access Protocol
SP	service pack
SSD	solid-state drive
SSH	Secure Shell
SSHA	Salted Secure Hash Algorithm
SSL	Secure Sockets Layer
SSO	single sign-on
SVGA	Super Video Graphics Array
TCP	Transmission Control Protocol
TFTP	Trivial File Transfer Protocol
TLS	Transport Layer Security
TOS	type of service
TTL	time to live
UAC	User Account Control
UDP	User Datagram Protocol
UID	user identifier
UNC	Universal Naming Convention
UPnP	Universal Plug and Play
URI	Uniform Resource Identifier
URL	Uniform Resource Locator
UTC	Coordinated Universal Time
UTF	UCS Transformation Format
VDEV	Virtual Device
VLAN	virtual LAN
VLL	Virtual LVI/LUN
WADL	Web Application Description Language
WAN	wide area network
WINS	Windows Internet Name Service
WORM	Write Once, Read Many
WS	workstation

WWN	World Wide Name
WWW	World Wide Web
XDR	External Data Representation
XFS	extended file system
XML	Extensible Markup Language

Index

Symbols

/etc/cifs/lmhosts (system file) 3-3
/etc/hosts (system file) 3-3, 14-3

Numerics

8.3 format
 relation to MS-DOS file names 8-4

A

access ACL 8-13
access control entry 8-5
access control list 6-2
access permission
 inheriting from parent folder 8-17
accessing
 CIFS shares 7-1
 NFS shares 17-1
 notes for file shares 19-2
ACE 8-5
 duplication check 8-34
 invalid 8-23, 8-35
ACE flag 8-28
ACE mask 8-28
ACE type 8-28
ACL 6-2, 8-5, 18-2
 adding in Windows Server 2008 11-2
 adding in Windows Server 2012 11-11
 default when ACL is not inherited 8-38
 evaluation 8-33
 for new file 8-22
 for new folder 8-23

 initial values, inheritance, and propagation 8-34
 mapping to HDI system file permissions 8-23
 maximum number that can be set 8-37
 related values 8-28
 setting and displaying from CIFS client 8-24
 share-level 9-10
 share-level in Windows Server 2008 11-3
 share-level in Windows Server 2012 11-11
ACL settings
 specifying from CIFS client 8-9
 specifying or viewing a file 8-11
 specifying or viewing folder 8-12
Active Directory
 registering IDs 4-17
Active Directory authentication 3-5
Active Directory domain controllers 2-4
ADAM G-2
 creating a schema file 4-12
 precautions when setting up an LDAP server 4-9
 setting index 4-15
 setting up LDAP server 4-12
adding
 NFS service principals C-3
 user ACLs or group ACLs 8-20, 8-44
 users 4-3
Advanced ACL 8-5, 18-2
 file system 6-2, 8-24
 migrating to Advanced ACL file system 8-37
anonymous user 14-5, 15-3
archive attribute
 notes on 8-49
authentication
 when user mapping is being used 5-7
authentication modes 3-4

in CIFS service configuration definition 3-4

B

backing up

most recent CIFS access logs 3-17

C

changing

NFS share settings 14-5

checking

file attributes from CIFS client 8-47

CIFS

multi-port configuration 2-7

CIFS access log

backing up most recent 3-17

CIFS administrator 6-5

CIFS client 2-2

HDI node on different network 2-7

HDI node on same subnetwork 2-6

platforms 11-1

user authentication 5-1

CIFS client users

managing 4-1

CIFS file share

changing share information 9-6

creating 3-10, 9-5

editing attributes 3-11

files and folders within 8-1

managing 3-10

viewing a list of 9-4

CIFS file share name

displaying 8-4

CIFS protocols 2-5

CIFS service

changing CIFS configuration definition 3-3

configuration definitions 3-3

overview of using 1-2

running by File Services Manager 3-1

setting authentication mode 3-4

setting user mapping 3-6

supported products 2-2

system configuration 2-1

CIFS service configuration definition

Setting the SMB protocol 3-7

CIFS shared directory

access method 7-2

note on Anti-Virus Enabler environment 7-10

notes on accessing 7-3

CIFS shares

accessing 7-1

Classic ACL 8-5, 18-2

file system 6-2, 8-7

configuring

KDC server C-3

network and system information 14-2

NFS environment for both CIFS and NFS 13-10

NFS environment for Kerberos authentication C-1

NFS environment for NFS service only 13-8

creating

NFS shares 14-5

CSV file format 4-4

D

DACL 8-5

date and time

file was accessed 8-51

file was created 8-52

file was modified 8-51

default ACL 8-13

definitions

CIFS service configuration 3-3

NFS service configuration definition 14-4

service configuration 3-3, 14-3

deleting

users 4-3

users by script 4-5

deleting registered IDs

using LDAP server 4-22

differences between Classic ACLs and Advanced

ACLs 8-6

directories

mounting root directory 17-4

mounting shared directories 17-2

names in NFS shares 18-2

NFS sharing 8-48

notes on modifying 19-3

discretionary access control list 8-5

disk capacity

displayed when multiple quotas are set 8-59

displaying 8-53, 18-4

displaying in accordance with disk usage 8-57

disk usage

displaying in accordance to disk capacity 8-57

- displaying
 - ACLs from CIFS client 8-24
 - CIFS file share name 8-4
 - disk capacity 8-53, 18-4
 - disk capacity in accordance with disk usage 8-57
 - disk capacity when multiple quotas are set 8-59
- distributing
 - keytab files C-18
- DNS
 - network configuration 2-7
- DNS domain 13-6
- DNS server 14-3
- domain
 - user management 4-8
- domain controllers
 - Active Directory 2-4
- duplication check
 - ACE 8-34

E

- editing
 - system files 14-3
- encoding 18-2
- extended attributes
 - in Windows 8-50

F

- FAQ A-22
- file
 - ACL for new file 8-22
 - NFS sharing 8-48
- file and directory names
 - about 8-2
 - supported characters 8-2
- file and folder
 - in shared directories in Windows 10 11-8
 - in shared directories in Windows 7 11-4
 - in shared directories in Windows 8 11-7
 - shared directories 11-2, 11-10
- file attribute 8-47
 - setting or checking from CIFS client 8-47
- file locking 17-7
- file owner
 - Advanced ACL 8-35
- file permission
 - mapping from Windows ACLs 8-23

- File Services Manager
 - running CIFS service 3-1
 - setup 14-2
 - setup procedure 3-2
 - specifying security flavor D-2
 - using to run NFS service 14-1
- file shares
 - cautionary notes 19-1
 - notes on accessing 19-2
- file system
 - Advanced ACL 6-2, 8-24
 - Classic ACL 6-2, 8-7
 - mounting and viewing 17-2
 - root ACL 8-27
- file timestamp
 - granting update permission 8-53
 - management method 8-52
 - update resolution 8-52
- file timestamp resolution 8-52
- files
 - attributes in HDI 18-2
 - names in NFS shares 18-2
- files and folders
 - in CIFS share 8-1
- folder
 - ACL for new folder 8-23
- format for group mapping file 4-5
- format for user registration file 4-5

G

- group ACLs
 - adding 8-20, 8-44
- group ID
 - registering using Active Directory 4-17
 - registering using LDAP server 4-21
- group mapping
 - registering or removing 4-4

H

- HDI
 - CIFS client on different network 2-7
 - CIFS client on same subnetwork 2-6
- HDI system
 - disk capacity displayed when multiple quotas are set 8-59
- home drive

setting 7-11

I

inheritance

ACL 8-34

inheriting

access permissions from parent folder 8-17

initial value

ACL 8-34

K

KDC server 13-8

before configuring C-3

configuring C-3

supported by NFS service 13-4

KDC server domain 13-6, 13-8

Kerberos

configuring NFS environment 13-7

configuring NFS environment for C-1

Kerberos authentication 13-8, 16-2

keytab file 13-8, 13-9, 13-11

distributing C-18

distribution destination C-18

retrieving (for HDI nodes) C-19

retrieving (for NFS client) C-19

keytab files

distributing and retrieving C-18

L

LDAP server 4-2

deleting registered IDs 4-22

registering IDs 4-20

registering information in 4-2

using ADAM to set up 4-9

using OpenLDAP to set up 4-9

using Sun Java System Directory Server to set up 4-10

LDAP server for user mapping

setting up 4-8

listing

users by script 4-5

local authentication 3-4

local user management 4-2

M

managing

CIFS client users 4-1, 4-2

CIFS file share 3-10

NFS client users 15-1

NFS shares 14-5

mapping

Windows ACLs to HDI system file permissions 8-23

message A-1

CIFS log A-3

syslog A-2

migrating

Advanced ACLs from Windows 8-39

to Advanced ACL file system 8-37

migrating user resources 6-1

MMC

CIFS share management from 9-4

linking to 9-1

managing open files from 9-9

notes on using 9-12

session management from 9-7

using in Windows 10 11-9

using in Windows 7 11-6

using in Windows 8 11-7

using in Windows Server 2008 11-3

using in Windows Server 2012 11-11

modifying

directories 19-3

mount command

example of execution 17-3, 17-5

mounting

file system 17-2

root directory 17-4

shared directories 17-2

MS-DOS file name

8.3 format 8-4

N

name resolution

host name 17-2

name resolution services 7-2

NetBIOS over TCP/IP 2-5

network configuration 2-5

CIFS client and HDI node on different subnetworks 2-7

- CIFS client and HDI node on same subnetwork 2-6
- with CIFS and CIFS services 13-6
- network configurations
 - CIFS on multiple ports 2-7
 - DNS used 2-7
 - when only NFS service is running 13-5
- network drive
 - using in Windows Server 2008 11-3
- network information
 - configuring 14-2
- NFS
 - configuring environment for Kerberos authentication C-1
 - sharing file or directory 8-48
- NFS client
 - mounting shared directories D-2
- NFS clients
 - managing users 15-1
 - supported by NFS service 13-2
 - user authentication 16-1
- NFS environment
 - configuring 13-7
- NFS protocol 12-2
- NFS Service
 - overview 12-1
- NFS service
 - network configuration 13-5
 - supported products 13-2
 - using File Services Manager to run 14-1
- NFS service configuration definition 14-4
- NFS service principals
 - adding C-3
- NFS shared directory
 - accessing D-3
- NFS shares
 - access method 17-2
 - accessing 17-1
 - creating and changing 14-5
 - editing attribute 14-6
 - files and directories 18-1
 - managing 14-5
- nfscacheflush command 15-3
- NFSv4
 - configuring NFS environment 13-7
- NFSv4 domain 13-6, 15-2
- NFSv4 domain name definition file 15-3
- NIS server 4-2
 - registering information in 4-2

- note on user authentication
 - Active Directory authentication 5-3
 - local authentication 5-2
 - NT domain authentication 5-2
- notes
 - editing CIFS share attributes 3-11
- NT domain authentication 3-4

O

- offline file
 - enabling in Windows 10 11-9
 - enabling in Windows 7 11-5
 - enabling in Windows 8 11-7
 - enabling in Windows Server 2008 11-3
- open file
 - closing 9-9
 - list of 9-9
- Open LDAP
 - setting the index directive 4-12
 - creating a schema file 4-11
- OpenLDAP
 - precautions when setting up LDAP server 4-9
 - setting up LDAP server 4-11

P

- principal 13-8
- propagation
 - ACL 8-34

Q

- quota
 - checking on CIFS client 8-55
 - disk capacity displayed when multiple ones are set 8-59
 - using in Windows 10 11-9
 - using in Windows 7 11-5
 - using in Windows 8 11-7
 - using in Windows Server 2008 11-3
 - using in Windows Server 2012 11-11
- quota information
 - setting up 3-12
- quota management functionality 8-53
- quotas
 - notes on 3-12

R

- read-only attribute
 - notes on 8-49
 - registering
 - group mapping 4-4
 - information in NIS server or LDAP server 4-2
 - user IDs and group IDs 4-17
 - users by script 4-5
 - registering IDs
 - using Active Directory 4-17
 - using LDAP server 4-20
 - removing
 - group mapping 4-4
 - resource migration
 - backup utility 6-7
 - before performing 6-2
 - checking CIFS log 6-9
 - creating backup file 6-8
 - creating file system and CIFS share 6-8
 - executing ACL settings 6-9
 - obtaining ACL information 6-7
 - obtaining file attribute 6-7
 - registering CIFS administrator 6-8
 - restoring backup file 6-8
 - retention period F-1
 - retrieving
 - keytab files C-18
 - root 17-4
 - root ACL
 - file system 8-27
- ## S
- SACL 8-5, 8-23, 8-35
 - script
 - CIFS group mapping 4-6
 - secondary KDC server
 - adding E-1
 - security flavor 16-2
 - specifying from File Services Manager D-2
 - service configuration
 - definitions 14-3
 - service configuration definition 3-3
 - session
 - closing 9-8
 - viewing a list of 9-7
 - setting
 - ACLs from CIFS client 8-24
 - file attributes from CIFS client 8-47
 - home drive 7-11
 - setting up
 - LDAP server 4-8
 - quota information 3-12
 - setting up LDAP server
 - ADAM 4-12
 - OpenLDAP 4-11
 - Sun Java System Directory Server 4-15
 - setup
 - File Services Manager 3-2, 14-2
 - of network and system information 3-2
 - setup procedure for File Services Manager 14-2
 - shared directories 17-1
 - shared directory 7-1
 - mounting from NFS clients D-2
 - specifying
 - ACL settings for file 8-11
 - ACL settings for folder 8-12
 - Advanced ACL settings from CIFS client 8-9
 - Sun Java System Directory Server
 - creating a schema file 4-15
 - precautions when setting up LDAP server 4-10
 - setting index 4-16
 - setting up LDAP server 4-15
 - supported by NFS service
 - ID mapping server 13-5
 - supported products
 - ID mapping server (LDAP server for user authentication) 13-5
 - ID mapping server (NFS service) 13-5
 - KDC server (Active Directory domain controller) 13-5
 - KDC server (UNIX) 13-4
 - NFS client 13-2
 - symbolic link 19-2
 - system access control list 8-5
 - system configuration
 - when CIFS service is used 2-1
 - system files
 - /etc/cifs/lmhosts 3-3
 - /etc/hosts 3-3, 14-3
 - editing 3-3, 14-3
 - system information
 - configuring 14-2

T

timestamp 8-51

U

UNIX (AUTH_SYS) authentication 16-2
UNIX permission
 Advanced ACL 8-35
user ACLs
 adding 8-20, 8-44
user authentication
 CIFS client 5-1
 for NFS clients 16-1
 Kerberos 16-2
 methods 16-2
 UNIX (AUTH_SYS) 16-2
user ID
 registering 4-18
 registering using LDAP server 4-21
user IDs and group IDs
 registering 4-17
user management
 methods for NFS client users 15-2
 with NFSv4 domain 15-2
user management methods (CIFS) 4-2
user mapping
 authentication while being used 5-7
 setup in CIFS service configuration definition 3-6
user registration
 notes on 4-8
users
 adding or deleting 4-3
 managing CIFS client users 4-2
 script for registering, deleting, listing 4-5
UTF-8 8-2, 18-2

V

viewing
 ACL settings for file 8-11
 ACL settings for folder 8-12
 file system 17-2

W

Windows
 logging on in Windows Server 2008 11-3

logging on in Windows Server 2012 11-11
notes on all supported types 11-2

Windows 10
 notes on 11-8
Windows 7
 notes on 11-4
Windows 8
 notes on 11-6
Windows server
 disk capacity when multiple disk quotas are set
 8-63
Windows Server 2008
 notes on 11-2
Windows Server 2012
 notes on 11-10
Windows Server 2016
 notes on 11-8
WORM file 18-3

X

XCOPY 6-5

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
HitachiVantara.com | community.HitachiVantara.com



Contact Information
USA: 1-800-466-0744
Global: 1-858-547-4526
HitachiVantara.com/contact