

ONCplus B.11.31.07.01 Release Notes

HP-UX 11i v3

HP Part Number: 5900-0306
Published: September 2009



© Copyright 2009 Hewlett-Packard Development Company, L.P

Legal Notices

Confidential computer software. Valid license required from Hewlett-Packard for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for Hewlett-Packard products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting additional warranty. Hewlett-Packard shall not be liable for technical or editorial errors or omissions contained herein.

Oracle is a registered US trademark of Oracle Corporation, Redwood City, California.

UNIX is a registered trademark of The Open Group.

Table of Contents

1 ONCplus Release Notes.....	5
Announcement.....	6
Availability of ONCplus as an ISU.....	6
NIS+ Obsolescence Notice.....	6
NIS Protocol Version 1 Obsolescence Notice.....	6
Overview.....	7
Related Information.....	7
What's New in ONCplus B.11.31.07.01.....	7
Fixes in ONCplus to date.....	7
Installation Information	10
System Requirements.....	10
Installing ONCplus B.11.31.07.01.....	10
Verifying ONCplus B.11.31.07.01 Installation.....	11
Reverting to a Previous Version of ONCplus.....	13
Features Introduced in ONCplus B.11.31.07.....	14
Features Introduced in ONCplus B.11.31.06.....	14
New NIS Features in ONCplus B.11.31.06.....	14
New NFS Features in ONCplus B.11.31.06.....	14
Features Introduced in ONCplus B.11.31.05.....	14
Features Introduced in ONCplus B.11.31.04.....	14
New NIS Features in ONCplus B.11.31.04.....	15
NGROUPS Expansion.....	15
Features Introduced in ONCplus B.11.31.03.....	15
New NFS Features in ONCplus B.11.31.03.....	16
Features Introduced in ONCplus B.11.31.02.....	16
New CacheFS Features in ONCplus B.11.31.02.....	16
New NIS Features in ONCplus B.11.31.02.....	17
Features Introduced in ONCplus B.11.31.01.....	17
Features Introduced in ONCplus B.11.31_LR	17
New NFS Features in ONCplus B.11.31_LR.....	17
New AutoFS Features in ONCplus B.11.31_LR.....	19
New CacheFS Features in ONCplus B.11.31_LR	20
New NIS Features in ONCplus B.11.31_LR.....	21
Software Availability in Native Languages.....	21

List of Tables

1-1	ONCplus Software Fixes.....	8
-----	-----------------------------	---

1 ONCplus Release Notes

This document contains the most recent product information pertaining to the Open Networking Computing (ONC) product, version B.11.31.07.01, which is supported on the HP-UX 11i v3 operating system. This document addresses the following topics:

- “Announcement” (page 6)
- “Overview” (page 7)
- “What’s New in ONCplus B.11.31.07.01” (page 7)
- “Fixes in ONCplus to date” (page 7)
- “Installation Information ” (page 10)
- “Features Introduced in ONCplus B.11.31.07” (page 14)
- “Features Introduced in ONCplus B.11.31.06” (page 14)
- “Features Introduced in ONCplus B.11.31.05” (page 14)
- “Features Introduced in ONCplus B.11.31.04” (page 14)
- “Features Introduced in ONCplus B.11.31.03” (page 15)
- “Features Introduced in ONCplus B.11.31.02” (page 16)
- “Features Introduced in ONCplus B.11.31.01” (page 17)
- “Features Introduced in ONCplus B.11.31_LR ” (page 17)
- “Software Availability in Native Languages” (page 21)

Announcement

This version of ONCplus (B.11.31.07.01) is supported on systems running the HP-UX 11i v3 operating system.

Availability of ONCplus as an ISU

Prior to the HP-UX 11i v3 release, ONC was delivered as a core product called NFS. Core products cannot be released independent of the operating system. Any changes (defect fixes, for example) to core products are only available to customers by installing HP-UX patches.

For HP-UX 11i v3, ONCplus is available as an Independent Software Unit (ISU). With the ONCplus ISU, both defect fixes and new features are made available to customers by installing newer versions of ONCplus.

The ONCplus ISU will be updated on an ongoing basis and will include the latest ONC updates and defect fixes. When looking for ONC updates, you are encouraged to review the content of the latest ONCplus ISU and consider updating to the most recent ONCplus ISU version.



NOTE: All ONCplus ISU versions are available at:

<http://software.hp.com>.

NIS+ Obsolescence Notice

Network Information Service Plus (NIS+) is a distributed database system that enables the maintenance of commonly used configuration information on a master server and propagates the information to all the hosts in the network. Starting with HP-UX 11i v3, NIS+ is no longer supported. Hewlett-Packard recommends that users migrate to LDAP.



NOTE: For information on how to migrate from NIS+ to LDAP, see the *NIS+ to LDAP Migration Guide* available at:

<http://docs.hp.com/en/J4269-90054/index.html>

NIS Protocol Version 1 Obsolescence Notice

Starting with ONCplus version B.11.31.02, NIS protocol version 1 (NISv1) is no longer supported on the NIS client. However, the NIS server, `ypserv`, will continue to support NISv1.

The NISv1 protocol obsolescence impacts an NIS client application that uses the NISv1 protocol in the following scenarios:

- The NIS client uses the `-v1` option of the `ypwhich` or the `ypset` command.
- An application directly communicates with the `ypbind` Remote Procedure Call (RPC) version 1.
- An application includes the `ypv1_prot.h` header file.



NOTE: Starting with ONCplus version B.11.31.02, the `ypwhich` or the `ypset` command communicates with `ypbind` RPC version 3. However, to specify the `ypbind` RPC version to be used, you must use the options provided with the `ypwhich` and `ypset` commands.

Overview

Open Network Computing (ONC) comprises core services that enable administrators to implement distributed applications in a heterogeneous distributed computing environment. It also includes tools to administer clients and servers.

ONC consists of the following components:

- Network File System
The Network File System (NFS) is a distributed filesystem that provides transparent access to files and directories that are shared on remote systems.
- AutoFS
AutoFS is a client-side service that enables automatic mounting and unmounting of filesystems.
- CacheFS
The Cache Filesystem (CacheFS) is a general purpose filesystem caching mechanism that improves the performance of client side applications when working with NFS servers. CacheFS client performance is improved by caching data to a fast local file system instead of going over the wire. Caching data results in reduced server and network load because the clients have already cached a copy of the data and send fewer requests to the server.
- Network Information Service
Network Information Service (NIS), previously called “Yellow Pages,” is a distributed database system that enables the maintenance of commonly used configuration information on a master server and propagates the information to all the hosts in the network.

Related Information

For more information about ONCplus, see the following documents, available at:

<http://docs.hp.com/en/netcom.html#NFS%20Services>

- *NFS Services Administrator's Guide*
- *NIS Administrator's Guide*
- *Managing NFS and KRPC Kernel Configurations in HP-UX 11i v3*
- *Introducing Network File Systems Version 4 on HP-UX 11i v3*

What's New in ONCplus B.11.31.07.01

ONCplus B.11.31.07.01 is a defect fix release and does not include any new features. All features introduced in previous ONCplus versions are supported in ONCplus B.11.31.07.01.

Fixes in ONCplus to date

This section lists the fixes for ONCplus and also specifies the version number of ONCplus in which the defect was fixed.



WARNING! *The fix for QXCR1000898226 will not be available when:

- Upgrading to ONCplus B.11.31.07.01 from ONCplus B.11.31.07
- Installing HP-UX 11i v3 September 2009 release on a system with ONCplus B.11.31.07

HP recommends that you to install ONCplus B.11.31.08 September 2009 or a later version to fix this problem.

Table 1-1 describes the problems that are fixed to date and incorporated into ONCplus.

Table 1-1 ONCplus Software Fixes

CR ID	Description	ONCplus Version
QXCR1000914413	Directory related operations on an NFS client with ONCplus B.11.31.06 or B.11.31.07 versions installed along with mounted file systems with read/write permissions that are greater than 8192 bytes in size may result in system panic or data corruption.	B.11.31.07.01
QXCR1000869504	System panics while renaming a file located on a local filesystem mounted with nfs4deleg stackable module.	B.11.31.07
QXCR1000833665	Performance improvement in rpc.statd to improve the crash recovery.	B.11.31.07
QXCR1000886408	Under certain circumstances, NFS client with tcp mount fails to connect to the server, even when the server is available on the network	B.11.31.07
QXCR1000890608	Numerous ypserv processes waiting for lockf() on map lock.	B.11.31.07
*QXCR1000898226	NFSv3 does not fully support large files	B.11.31.07
QXCR1000893901	Support for chained mblks to avoid any data corruption during a read operation.	B.11.31.07
QXCR1000857730	On NFS mounts with tcp and noac options, setting nfs3_bsize tunable to a value greater than or equal to 64KB, results in file sizes more than its actual size at the time of their creation or modification.	B.11.31.07
QXCR1000855994	Under certain circumstances, enabling the nfs2_srv_read_copyavoid or nfs3_srv_read_copyavoid tunable may cause the following problems: - Memory corruption with the message block getting corrupted - Data corruption during READ, when the nfs3_bsize tunable value is greater than 64 KB	B.11.31.07
QXCR1000797858	NFSv3 and NFSv4 set the transfer size incorrectly.	B.11.31.06
QXCR1000828306	Compatibility issue with HP-UX 11i v2 xdr_callmsg(), xdr_callhdr(), and xdr_replymsg().	B.11.31.06
QXCR1000864699	NFS client data corruption due to out-of-order write replies.	B.11.31.06
QXCR1000852442	Under certain circumstances, there is a possibility of a memory leak in NFS.	B.11.31.06
QXCR1000886293	System panics in clnt_clts_dispatch_notify().	B.11.31.06
QXCR1000795647	ONC 2.5 functionality for the NFSv2 and NFSv3 client.	B.11.31.06
QXCR1000808831	IPv6 support for RPC commands and daemons Phase II	B.11.31.06
QXCR1000589225	Added support for the ip6.arpa domain to the getnameinfo(3n) command.	B.11.31.05
QXCR1000847469	Under certain circumstances, ONCplus B.11.31.03 and B.11.31.04 may cause a panic.	B.11.31.05
QXCR1000590736	rpc.yupdated, ypxfrd, and keyserv are not killed when the NIS Master Server is enabled and disabled by smh.	B.11.31.04
QXCR1000593073	Documentation issues in /etc/rc.config.d/namesvrs.	B.11.31.04
QXCR1000752996	Support for NGROUPS expansion in ONC.	B.11.31.04
QXCR1000756355	IPv6 support for RPC commands and daemons Phase I	B.11.31.04
QXCR1000774733	Optimize direct I/O for NFSv2 and NFSv3 reads and writes.	B.11.31.04
QXCR1000779532	Remove nfs4_nrnode and nfs4_nacache private tunables from 11i v3.	B.11.31.04
QXCR1000785512	An NFSv4 mount overwrites the mount error returned by the server with ENOENT.	B.11.31.04
QXCR1000790796	NUMA support changes for CacheFS.	B.11.31.04

Table 1-1 ONCplus Software Fixes (continued)

CR ID	Description	ONCplus Version
QXCR1000792753	Spinlock contention in checkexport/exi_rele.	B.11.31.04
QXCR1000795425	Serviceguard NFS failover fails in the presence of the "fsid=" option.	B.11.31.04
QXCR1000810049	Remove the NFSv4.1 server stubs for referrals.	B.11.31.04
QXCR1000816572	The xdr_hyper() and xdr_u_hyper() functions incorrectly process data.	B.11.31.04
QXCR1000592851	Serviceguard NFS Toolkit returns a stale file handle on a server failover.	B.11.31.03
QXCR1000586579	Added support for NFSv4 referrals, cross-mounts, and local access delegation stackable module.	B.11.31.03
QXCR1000742446	A multi-threaded process reading an NFS directory hangs and cannot be killed.	B.11.31.03
QXCR1000745905	An NFS close() returns the error message "close: operation would block".	B.11.31.03
QXCR1000743541	An NFS file is corrupted under heavy system load.	B.11.31.03
QXCR1000759153	AutoFS fails to mount with NFSv4 after an update from HP-UX 11i v2 to HP-UX 11i v3.	B.11.31.03
QXCR1000769108	A system runs out of memory and hangs when NFS and NFS file locking are heavily used.	B.11.31.03
QXCR1000763503	An AutoFS UDP NFS mount fails with the error message "Couldn't bind to reserved port".	B.11.31.03
QXCR1000764069	Added support for SecureNFS in a Serviceguard environment.	B.11.31.03
QXCR1000764257	An NFS server panics when using a public filehandle for a lookup.	B.11.31.03
QXCR1000771618	There is a potential race condition on an NFSv4 client using direct I/O that can result in a hang.	B.11.31.03
QXCR1000770472	NFS REaddirPLUS returns no entries when rsize is set too low.	B.11.31.03
QXCR1000584936	Performance improvement in NFS client WRITE path.	B.11.31.02
QXCR1000585229	Added support for ACL's and logging in CacheFS.	B.11.31.02
QXCR1000587643	Added an interface for assigning a port number for the rpc.pcnfsd daemon.	B.11.31.02
QXCR1000590350	The nfs4bcd daemon is not terminated by "nfs.client stop" after enabling NFS client in smh.	B.11.31.02
QXCR1000590569	The share_nfs(1M) manpage must warn about the limitations of the -access option.	B.11.31.02
QXCR1000590702	An NFS mount fails if the shared hostname is longer than eight characters.	B.11.31.02
QXCR1000590739	The automountd daemon is not terminated by "autofs stop" after enabling automounter in smh.	B.11.31.02
QXCR1000591016	CacheFS does not support locking a region of a file beyond a 2 TB offset.	B.11.31.02
QXCR1000591305	The readdirplus transfer size should be greater than 8K.	B.11.31.02
QXCR1000591307	Added support to disable readdirplus functionality from the NFS mount command.	B.11.31.02
QXCR1000591590	An NFSv4 server panics in ACL code.	B.11.31.02
QXCR1000591875	Added functionality for the NIS 2.3 client.	B.11.31.02
QXCR1000592166	An AutoFS client using NFSv4 cannot automount an NFS server which does not support NFSv4.	B.11.31.02
QXCR1000592981	The automountd daemon hangs while halting a package.	B.11.31.02

Table 1-1 ONCplus Software Fixes (continued)

CR ID	Description	ONCplus Version
QXCR1000593322	Performance improvement in NFS Server WRITE path.	B.11.31.02
QXCR1000731755	The automountd daemon does not work correctly with long hostnames.	B.11.31.01
QXCR1000732349	The rpc.pcnfsd and rpc.rexd daemons do not support long usernames.	B.11.31.01
QXCR1000735912	System panics with heavy load on a CacheFS mount point; CacheFS unmount threads may hang.	B.11.31.01
QXCR1000581253	Problem with rpc.yppasswd.	B.11.31.01
QXCR1000582801	A user belonging to more than 16 groups cannot access files over NFS.	B.11.31.01
QXCR1000582803	Files created on an NFSv4 server have an incorrect group id.	B.11.31.01
QXCR1000582821	The output for the nfsstat command used with -m option is incorrect when using replicated servers or long hostnames.	B.11.31.01
QXCR1000583678	CacheFS panics when issuing simultaneous mount requests with a full front file system.	B.11.31.01
QXCR1000585325	The nfsstat.h header file incorrectly defines MAX_RFS4_PROC_NO as 40.	B.11.31.01
QXCR1000585353	An NFS server panics due to a data page fault.	B.11.31.01
QXCR1000585618	Added a new private kctune parameter "klm_log_level" to enable KLM (kernel lock manager) logging.	B.11.31.01
QXCR1000585874	The nfsstat command does not display access, close, and commit values for NFSv4 when used with the -c, or -s options.	B.11.31.01
QXCR1000585903	Using kctune to change NFSv4 parameters causes a panic.	B.11.31.01
QXCR1000585959	Users cannot access an NFS mounted file system in PAM enabled Kerberos systems.	B.11.31.01
QXCR1000586940	The xdr_setpos function in libnsl.1 does not work properly.	B.11.31.01
QXCR1000587604	Writing from an NFSv3 client to a Red Hat Linux server fails with an "RPC: Server can't decode arguments" error message.	B.11.31.01
QXCR1000588479	The rpcbind daemon does not receive IPv6 multicast.	B.11.31.01

For more information on these defects and how they were fixed, see:

<http://www2.itrc.hp.com>.

Installation Information

This section elaborates on the requirements and the procedures for installing ONCplus B.11.31.07.01. It also includes information on verifying the installation of ONCplus B.11.31.07.01.

System Requirements

The following requirements are necessary for installing ONCplus B.11.31.07.01 :

- Hardware Requirement: Itanium / PA-RISC
- Software Requirement: HP-UX 11i v3
- Free Disk Space: 65 MB

Installing ONCplus B.11.31.07.01



IMPORTANT: Back up your system before installing the product.

To install ONCplus B.11.31.07.01 on your system:

1. Login as root.
2. To check which version of ONCplus is currently installed on your system, run the following command:

```
swlist | grep ONCplus
```

The output will be similar to one of the following nine lines:

```
ONCplus      B.11.31.07      ONC+ 2.3
ONCplus      B.11.31.06.01  ONC+ 2.3
ONCplus      B.11.31.06      ONC+ 2.3
ONCplus      B.11.31.05      ONC+ 2.3
ONCplus      B.11.31.04      ONC+ 2.3
ONCplus      B.11.31.03      ONC+ 2.3
ONCplus      B.11.31.02      ONC+ 2.3
ONCplus      B.11.31.01      ONC+ 2.3
ONCplus      B.11.31         ONC+ 2.3
```

3. After you download the ONCplus_B.11.31.07.01.depot file, move it to the /tmp directory:

```
/tmp/ONCplus_B.11.31.07.01.depot
```

4. Verify that the file has downloaded correctly using the swlist command as follows:



NOTE: You must specify the full path name of the source depot when you use swlist and swinstall commands.

```
swlist -d @ /tmp/ONCplus_B.11.31.07.01.depot
```

If ONCplus is downloaded correctly, the output will include:

```
#
# Bundle(s) :
#
      ONCplus                B.11.31.07.01      ONC+ 2.3
#
# Product(s) not contained in a Bundle:
#
      PHCO_38048            1.0                libc cumulative patch
```

5. Run the following command to install the product on a stand-alone system:

```
swinstall -x autoreboot=true -s \
/tmp/ONCplus_B.11.31.07.01.depot ONCplus
```



NOTE: ONCplus includes kernel filesets. As a result, installing the product using the swinstall command will require the system to be restarted after the installation is complete.

Verifying ONCplus B.11.31.07.01 Installation

To verify the ONCplus installation:

1. Run the following command to verify the software installation:

```
swverify ONCplus
```

If ONCplus B.11.31.07.01 is successfully installed, the following message is displayed:

```
Verification succeeded
```

2. To check which version of ONCplus is installed on your system, issue the following command:

```
swlist | grep ONCplus
```

The output will be similar to the following:

```
ONCplus      B.11.31.07.01  ONC+ 2.3
```

Reverting to a Previous Version of ONCplus

You cannot remove the current version of ONCplus with the `swremove` command. However, you can revert to the previous version of ONCplus. The following section describes how to revert to the previous version of ONCplus.



WARNING! Directory related operations on an NFS client with ONCplus B.11.31.06 or B.11.31.07 versions installed along with mounted file systems with read/write permissions that are greater than 8192 bytes in size may result in system panic or data corruption. If your system has ONCplus B.11.31.06 or B.11.31.07 version installed, HP recommends that you install ONCplus B.11.31.07.01 version.

WARNING! Under certain circumstances, enabling the `nfs2_srv_read_copyavoid` or `nfs3_srv_read_copyavoid` tunable may cause the following problems:

- Memory corruption with the message block getting corrupted
- Data corruption during READ, when the `nfs3_bsize` tunable value is greater than 64 KB

If your system has ONCplus B.11.31.02 or B.11.31.03 or B.11.31.04 or B.11.31.05 or B.11.31.06 version installed, HP recommends that you install ONCplus B.11.31.07 or a later version

WARNING! Under certain circumstances, ONCplus B.11.31.03 and B.11.31.04 may cause a panic situation (QXCR1000847469). If your system has either ONCplus B.11.31.03 or B.11.31.04 installed, HP recommends that you install ONCplus B.11.31.06.

WARNING! ONCplus B.11.31.02 may cause a boot panic situation if you are installing on an HP 9000 rp3410 or rp3440 system with the September 2007 release of HP-UX already installed. If your system has this configuration, Hewlett-Packard recommends that you install the March 2008 release of HP-UX, which includes ONCplus B.11.31.02, or install ONCplus B.11.31.06. You can use the "model" command to print the hardware model information. The output of the "model" command on an HP 9000 rp3440 system is:

```
9000/800/rp3440
```

To determine if you have the September 2007 release of HP-UX installed, issue the command:

```
swlist -l bundle HPUX11i-*
```

The output includes either "B.11.31" or "B.11.31.yymm", where yymm identifies the year and month of the HP-UX release that is installed. For example, if you have the HP-UX Mission Critical Operating Environment and the September 2007 release of HP-UX installed, the output would include:

```
HPUX11i-OE-MC B.11.31.0709 HP-UX Mission Critical Operating Environment
```



NOTE: ONCplus B.11.31.05 includes the fix for QXCR1000589225 which requires libc patch PHCO_38048 to be installed in the following circumstances:

- the "retrans" option is set in the `/etc/resolv.conf` file
- there is a "dns" entry in the `/etc/nsswitch.conf` file (i.e. the `ip6.arpa` domain is used to perform reverse IPv6 lookups) The libc patch is included with ONCplus beginning with version B.11.31.06 and will not require a separate patch installation.

To revert to a previous version of ONCplus, follow these steps:

1. Download the desired ONCplus depot from the ONCplus web page <http://software.hp.com>, or from the media you originally used to install your system.
2. Login as `root`.
3. Copy the downloaded ONCplus version of the depot (B.11.31_LR, B.11.31.01, B.11.31.02, B.11.31.03, B.11.31.04, B.11.31.05, B.11.31.06, B.11.31.07) to `/tmp/ONCplus_B.11.31.depot`.

4. To revert to the previous version of ONCplus enter the following command:

```
swinstall -x allow_downdate=true -x autoreboot=true \  
-s /tmp/ONCplus_B.11.31.depot ONCplus
```



NOTE: ONCplus contains kernel filesets. Thus swinstalling the product results in a system reboot after the installation is complete.

Features Introduced in ONCplus B.11.31.07

ONCplus B.11.31.07 is a defect fix release and does not include any new features. All features introduced in previous ONCplus versions are supported in ONCplus B.11.31.07.

For information on the defects fixed, see “Fixes in ONCplus to date” (page 7).

Features Introduced in ONCplus B.11.31.06

ONCplus B.11.31.06 includes both defect fixes and new features. This version of ONCplus supports all features included in previous ONCplus versions. In addition, the following new features are introduced in ONCplus B.11.31.06:

New NIS Features in ONCplus B.11.31.06

NIS introduces the following feature:

- **IPv6 support for RPC commands and daemons Phase II**

ONCplus B.11.31.06 delivers IPv6 support for `rpc.rquotad`, `rpc.rstatd`, `rpc.sprayd`, `rup`, and `spray` commands and daemons.

New NFS Features in ONCplus B.11.31.06

NFS introduces the following feature:

- **ONC 2.5 functionality for the NFSv2 and NFSv3 client**

Porting of ONC2.5 Solaris code base to the current HP-UX NFS v2 and NFS v3 client.

Features Introduced in ONCplus B.11.31.05

ONCplus B.11.31.05 is a defect fix release and does not include any new features. All features introduced in previous ONCplus versions are supported in ONCplus B.11.31.05.

For information on the defects fixed, see “Fixes in ONCplus to date” (page 7).

Features Introduced in ONCplus B.11.31.04

ONCplus B.11.31.04 includes both defect fixes and new features. All features introduced in previous ONCplus versions are supported in ONCplus B.11.31.04. The following features were introduced in ONCplus B.11.31.04:

New NIS Features in ONCplus B.11.31.04

NIS introduces the following features:

- **IPv6 support for RPC commands and daemons Phase I**

ONCplus B.11.31.04 delivers IPv6 support for `rpc.rexd`, `rpc.rwalld`, `rpc.rusersd`, `rusers`, and `rwall` commands and daemons. Other RPC commands and daemons (`rpc.statd` and `rpc.rquotad`) will support IPv6 in a future release of ONCplus.

NGROUPS Expansion

The NGROUPS expansion enhancement to HP-UX 11i v3 Update 3 changes the maximum number of groups for a user or process to be a tunable parameter. This allows customers to implement models for file access and protection that were not possible under the previous limit of 20 groups per user. Before deploying such models, customers must update the HP-UX kernel and utilities, and ensure that their local applications are compatible with the expanded group limit. ONCplus has the following limitations regarding the number of groups per user:

- The records, which can be passwords, groups, hosts, `netid`, etc., in the NIS database are limited to 1024 characters. The use of extremely long group names can cause the 1024 character limit to be exceeded even with just 20 groups, and in that case, such long records will not be included in the NIS database.
- NIS builds a database containing information about user group membership, called `netid`. NIS will include only the first 20 supplementary user groups in this database. This is usually used in conjunction with ONC RPC, where only the first 16 groups are significant. NIS `netid` maps will support only 20 user supplementary groups for backward compatibility.

For more information on the NGROUPS expansion enhancement to HP-UX 11i v3, refer to the NGROUPS Expansion: Guidelines for Deployment white paper at: <http://docs.hp.com/en/oshpux11iv3.html#White%20Papers>.

Features Introduced in ONCplus B.11.31.03

ONCplus B.11.31.03 includes both defect fixes and new features. All fixes and features introduced in previous ONCplus versions are supported in ONCplus B.11.31.03. The following features were introduced in ONCplus B.11.31.03:

New NFS Features in ONCplus B.11.31.03

NFS introduces the following features for NFSv4:

- **File Delegation with Local Access**

NFSv4 clients support delegation on HP-UX 11i v3. However, until ONCplus B.11.31.03, NFSv4 servers supported delegation with the caveat that no local file access will occur on any delegated file. For example, if the server grants a delegation to the file “/a/b/foo” then any local users on the server need to avoid accessing file “foo” while the delegation is in effect. If both local and remote users modify the delegated file, then the data in “foo” could become corrupted.

As of ONCplus B.11.31.03, when using the File Delegation, both local and remote file users can modify the delegated file.

- **Cross Mount Traversal**

The NFSv4 protocol allows clients to seamlessly traverse the servers shared directories and cross the physical file system boundaries on the server without having to explicitly mount each shared file system independently. For example, if the server is sharing the two file systems “/” and “/a/b” respectively, the client, after mounting the root file system of the server, can traverse the file system “/a/b” on the server without mounting the file system explicitly.

- **Referrals and Multi-server Namespace**

The Cross Mount Traversal feature allows an NFSv4 client to traverse the servers shared directories and seamlessly cross the physical file system boundaries on the server. The Referrals feature allows an NFSv4 client to traverse shared directories and seamlessly cross the physical file systems located on different servers. In other words, a referral defines a way for the NFSv4 server to direct an NFSv4 client to a file system which resides on a different server. The combination of cross mounts and referrals can be used to construct a global namespace.

For additional details on these new NFS features, please see the *Introducing Network File systems Version 4 on HP-UX 11i v3* white paper at: <http://docs.hp.com/en/netcom.html#NFS%20Services>

Features Introduced in ONCplus B.11.31.02

ONCplus B.11.31.02 includes both defect fixes and new features. All fixes and features introduced in previous ONCplus versions are supported in ONCplus B.11.31.02. The following features were introduced in ONCplus B.11.31.02:

New CacheFS Features in ONCplus B.11.31.02

CacheFS introduces the following features:

- **Support for ACLs**

An access control list (ACL) offers stronger file security by enabling the owner of the file to define file permissions for specific users and groups. HP-UX supports two types of ACLs: `HPUX_ACLS` and `SYSV_ACLS`. `HPUX_ACLS` are non-POSIX compliant. `SYSV_ACLS` are POSIX compliant. This version of CacheFS supports caching for only `SYSV_ACLS`. Thus, this version of CacheFS on HP-UX supports ACLs with VxFS and NFS and not with HFS.

- **Support for Logging**

A new command, `cachefslog` enables or disables logging for a CacheFS mount-point. If logging functionality is enabled, details about the operations performed on the CacheFS mount-point are stored in a logfile. This logfile contains information on all the CacheFS mount points using the same cache directory. Use the `cachefswssize` command to display the amount of space taken by each of the filesystems in the same cache and the total size

occupied by the cache directory (also known as the working set size). This command uses the logfile created by the `cachefslog` command to display the information. The `cachefswsize` command, used with the `-a` option, displays the information in ASCII format.

New NIS Features in ONCplus B.11.31.02

NIS introduces the following features:

- **Support for IPv6 Protocol**

The NIS clients and servers are now IPv6 enabled. Therefore, you can set up an NIS master server, an NIS slave server, or an NIS client that can be identified using an IPv6 address. The IPv6 information is stored in `ipnodes`, which are supersets of hosts that act as the new databases for IPv6 information.

- **Support for Resolving Map Nicknames**

NIS supports the creation of nicknames for maps. You can create or update the nicknames associated with the maps. The `ypcat` and `ypmatch` commands use the `/var/yp/nicknames` file to resolve or modify nicknames.

- **Support for NIS ypbind v3 Protocol**

The NIS client supports version 3 of the `ypbind` protocol. Version 1 of the `ypbind` protocol is obsolete and any request from an NIS client to `ypbind` version 1 is rejected.

- **Reduced Usage of Reserved Ports**

Reserved ports are the ports from 0 to 1024. Only root users can bind to these ports. In previous releases, NIS commands attempted to bind to reserved ports by default. If there are numerous client requests, all the reserved ports can be consumed. This version of NIS enables binding to reserved ports for select commands or daemons when accessing secure maps which results in reduced usage of reserved ports by NIS. This change does not compromise performance or security.

Features Introduced in ONCplus B.11.31.01

ONCplus B.11.31.01 is a defect fix release and does not include any new features. All features included in ONCplus B.11.31_LR continue to be supported in ONCplus B.11.31.01.

For information on the defects fixed, see “Fixes in ONCplus to date” (page 7).

Features Introduced in ONCplus B.11.31_LR

The following features were introduced in ONCplus B.11.31_LR:

New NFS Features in ONCplus B.11.31_LR

- **NFS Version 4 Protocol (NFSv4)**

NFSv4 is an IETF standard protocol that provides the following features:

- **COMPOUND Procedure**

In NFSv4, related RPC requests are grouped into a single RPC procedure known as the COMPOUND procedure.

The COMPOUND procedure decreases transport and security overhead because of fewer over-the-wire trips between the client and the server. This feature is transparent to the user.

- **Delegation**

In NFSv4, the server can delegate certain responsibilities such as, OPEN, CLOSE, LOCK, LOCKU, READ, and WRITE to the client. Delegation enables a client to locally service

operations without immediate interaction with the server. Delegations can be revoked by the server. When another client requests access to the same file, the server revokes the delegation from the first client and grants it to the second client.



NOTE: Delegations are disabled by default. If delegations are enabled, they are only supported by applications that access the delegated files remotely via NFS clients. Allowing local access and enabling delegation on a file can corrupt the file.

— **ACLs**

An access control list (ACL) offers stronger file security by enabling the owner of the file to define file permissions for the file owner, the group, and other specific users and groups.

— **String Identifiers**

In NFSv4, exchange of user ID (UID) and group ID (GID) information between the client and server is in the form of strings. The `nfsmapid` daemon maps UID and GID values from integer to string and string to integer.

— **Locks**

Locking support is integrated with the NFSv4 protocol. NFSv4 introduces leases for lock management. When a server grants a lock to control the state of a file for a specific period of time, it is known as a lease.

— **Single Protocol**

The MOUNT, Network Lock Manager (NLM), and the Network Status Monitor (NSM) protocols are merged into the NFS protocol. Merging these protocols into a single protocol enables easy configuration of firewalls. For more information, see *NFS Services Administrator's Guide*.



NOTE: The default NFS protocol version on HP-UX 11i v3 is 3. For information on how to configure NFSv4 as the default version, see *NFS Services Administrator's Guide* (B1031-90067).

• **Support for WebNFS**

WebNFS is an extension of the NFS protocol. It enables easy access to files across the Internet. Filesystems on the Internet can appear to a user as a local filesystem. WebNFS works through firewalls and enables applications that run on heterogeneous operating systems to access shared files.

• **Secure NFS**

With Secure NFS security can be implemented at the Remote Procedure Call (RPC) level. This standard authentication system is known as Secure RPC. When NFS uses the facilities provided by Secure RPC, it is known as Secure NFS.

RPC authentication enables you to use a variety of authentication systems, such as DH, UNIX, and KERB, and is not dependent on the version of NFS used.

• **Client Failover**

When an NFS server fails, the client accessing the shared files on that server can no longer access the shared files. If client failover is enabled, the client is automatically switched to an alternate server, which is a replica of the server that failed. The client continues to access shared files, without being aware of the switch.



NOTE: Client failover support is limited to read-only mounts or static filesystems that are not modified often.

- **Enhanced NFS logging**

NFS server logging enables an NFS server to provide a record of file operations that are performed on its filesystems.

New AutoFS Features in ONCplus B.11.31_LR

AutoFS supports the following features:

- **On-Demand Mounting of Hierarchical Filesystems**

In earlier versions, AutoFS mounted an entire set of filesystems if they were hierarchically related. However, in HP-UX 11i v3, AutoFS mounts only those filesystems that users access. Other filesystems that are hierarchically related to these filesystems are mounted only when they are requested. Hierarchical filesystem mounting prevents unnecessary mounting and unmounting, and improves performance.

- **Browsability**

AutoFS enables a user to browse the potential mount-points for indirect maps without actually mounting each filesystem.

- **Concurrent mount or unmount**

AutoFS performs concurrent mounts and unmounts using the multithreaded automountd daemon. Concurrent mounting or unmounting prevents services from hanging if a server is unavailable.

- **Reliable NFS ping**

AutoFS supports a `-retry=n` mount option for an NFS map entry to configure the ping timeout value. The default timeout value is 10 seconds. The ping command can impose a load on the network if used during normal operations or from automated scripts.

- **NFS loopback mount**

By default, AutoFS uses the Local File Systems (LOFS) mounts for locally mounted filesystems. AutoFS provides an option to allow loopback NFS mounts for the local mount. Use this option in high availability NFS environments.

- **Client-side Failover Support**

If the current server goes down, AutoFS enables a mounted NFS read-only filesystem to transparently switch over to an alternate server.

- **Backend Support**

AutoFS supports the storage and distribution of AutoFS maps in the following:

- Files
- Network Information Service (NIS)
- Lightweight Data Access Protocol (LDAP)

- **Filesystem Support**

AutoFS supports the automatic mounting and unmounting of the following filesystems:

- NFS (All versions including NFSv4)
- CacheFS
- HFS
- VxFS

- CIFS
- AutoFS
- **Secure NFS Support**
If the NFS client supports mounting of secure directories, AutoFS supports Secure NFS filesystems.
- **IPv6 Support**
AutoFS provides support for mounting filesystems over IPv6 transports.

New CacheFS Features in ONCplus B.11.31_LR

CacheFS supports the following features:

- **Complete Binary Caching**
CacheFS is commonly used to manage application binaries. This feature enables you to cache a complete binary file in the local cache. To force a client to cache a complete copy of the accessed binary file, use the `rpages` mount option. For more information on the `rpages` mount option, see the *NFS Services Administrator's Guide*.
- **Cache Pre-Loading**
This feature enables you to pre-load or pack specific files and directories in the cache. Packing files and directories enables you to have greater control over the cache contents because it ensures that the specified files are always present in the cache. Use the `cachefspack` command to pack specific files and directories.
- **Cache Administering**
This feature enables you to perform administrative tasks, such as creation and deletion of a cache directory. It also enables you to update the resource parameters of the specified cache directory. In addition, you can view the contents and statistics of the cache.
- **CacheFS Write Modes**
CacheFS supports two write modes, `write-around` and `non-shared`. In the `write-around` mode, writes are made to the back filesystem. `Write-around` mode is the default write mode. In the `non-shared` mode, writes are made both to the front and the back filesystems.
- **Cache Consistency Checking**
CacheFS checks the files that are stored in the cache to ensure that the files are kept up to date. With this version of CacheFS, the default behavior is that the consistency checking used can impact CacheFS performance. Following are mount options that can be used to change the type of consistency checking performed by CacheFS:
 - `noconst` Disable consistency checking.
 - `demandconst` Consistency checking is performed on demand.
 - `weakconst` Consistency check used to verify cache consistency with the NFS client's copy of the attributes. `weakconst` is the consistency checking level similar to the default behavior in HP-UX 11i v2.



NOTE: Consistency is not checked at file open time.

- **Switching Mount Options**

You can switch between mount options without deleting or rebuilding the cache. For instance, you can switch from `default` to `non-shared`, or from `noconst` to `demandconst` mount options without recreating the cache.

- **Support for Large Files and Large Filesystems**

CacheFS supports the maximum file and filesystem sizes supported by both the underlying front filesystem in which the cache resides and the back filesystem.

CacheFS data structures are 64-bit compliant.



NOTE: CacheFS does not support NFSv4 filesystems.

New NIS Features in ONCplus B.11.31_LR

NIS supports the following features:

- **Shadow Mode**

Starting with HP-UX 11i v3, NIS supports shadow password mode. This enables the NIS subsystem to recognize a system in shadow mode and to build, store, and retrieve the password information accordingly.



NOTE: The encrypted password information used for creating NIS password maps is present in the shadow file and is visible in the `passwd` maps (`passwd.byname` and `passwd.byuid`).

- **DNS Forwarding Mode**

The DNS Forwarding Mode enables the NIS server to contact the DNS server to service hosts or ipnodes requests. NIS can automatically fetch the information from DNS. This operation is transparent to the user or application and can be achieved without any change to the switch configuration.

- **Multi-homed Node**

NIS supports systems with more than one network address. This enables an NIS client to obtain the closest address of the host when it looks for a specific host name.

- **IPv6 Data Support**

NIS retrieves the IP addresses corresponding to the host name and identifies the IPv6 address formats. IPv6 data support is provided with ONCplus, but NIS is not IPv6 enabled.

- **Ipnodes Support**

Ipnodes is the new map that stores IPv6 information. Ipnodes acts as the new `hosts` database for IPv6 information. NIS provides support for ipnodes using the `/etc/nsswitch.conf` file.

- **Alternate Directory for passwd File**

Users can now specify a directory other than the default `/etc` directory where NIS must look for the `passwd` file. This feature enhances the security of the system, as information regarding this alternate directory is known only to the administrator.

Software Availability in Native Languages

The ONCplus software is only available in English.