

ONCplus B.11.31.17 Release Notes

HP-UX 11i v3

HP Part Number: 5900-3294
Published: September 2013
Edition: 1



© Copyright 2009, 2013 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.

Contents

1 ONCplus Overview	4
Overview.....	4
Availability of ONCplus as an ISU.....	4
NIS+ Obsolescence Notice.....	5
NIS Protocol Version 1 Obsolescence Notice.....	5
Related Information.....	5
2 ONCplus B.11.31.17	6
Announcement.....	6
What's New in ONCplus B.11.31.17.....	6
Enhancements and Fixes in ONCplus to date.....	6
Installation Information.....	14
System Requirements.....	14
Installing ONCplus B.11.31.17.....	15
Verifying ONCplus B.11.31.17 Installation.....	16
Reverting to a Earlier Version of ONCplus.....	16
3 Features Introduced in previous versions of ONCplus	24
Features Introduced in ONCplus B.11.31.16.....	24
Features Introduced in ONCplus B.11.31.15.....	24
Features Introduced in ONCplus B.11.31.14.....	24
Features Introduced in ONCplus B.11.31.13.....	25
Features Introduced in ONCplus B.11.31.12.....	25
Features Introduced in ONCplus B.11.31.11.....	25
Features Introduced in ONCplus B.11.31.10.....	25
Features Introduced in ONCplus B.11.31.09.02.....	26
Features Introduced in ONCplus B.11.31.09.01.....	26
Features Introduced in ONCplus B.11.31.09.....	26
New NFS Feature in ONCplus B.11.31.09.....	26
Features Introduced in ONCplus B.11.31.08.....	26
New NFS Features in ONCplus B.11.31.08.....	27
Features Introduced in ONCplus B.11.31.07.01.....	27
Features Introduced in ONCplus B.11.31.07.....	27
Features Introduced in ONCplus B.11.31.06.....	27
New NIS Features in ONCplus B.11.31.06.....	28
New NFS Features in ONCplus B.11.31.06.....	28
Features Introduced in ONCplus B.11.31.05.....	28
Features Introduced in ONCplus B.11.31.04.....	28
New NIS Features in ONCplus B.11.31.04.....	28
NGROUPS Expansion.....	28
Features Introduced in ONCplus B.11.31.03.....	29
New NFS Features in ONCplus B.11.31.03.....	29
Features Introduced in ONCplus B.11.31.02.....	29
New CacheFS Features in ONCplus B.11.31.02.....	29
New NIS Features in ONCplus B.11.31.02.....	30
Features Introduced in ONCplus B.11.31.01.....	30
Features Introduced in ONCplus B.11.31_LR.....	30
New NFS Features in ONCplus B.11.31_LR.....	31
New AutoFS Features in ONCplus B.11.31_LR.....	32
New CacheFS Features in ONCplus B.11.31_LR.....	33
New NIS Features in ONCplus B.11.31_LR.....	34
Software Availability in Native Languages.....	35

1 ONCplus Overview

This chapter addresses the following topics:

- “Overview” (page 4)
- “Availability of ONCplus as an ISU” (page 4)
- “NIS+ Obsolescence Notice” (page 5)
- “NIS Protocol Version 1 Obsolescence Notice” (page 5)
- “Related Information” (page 5)

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.

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. HP recommends that users migrate to LDAP.

NOTE: For information on how to migrate from NIS+ to LDAP, see the *NIS+ to LDAP Migration Guide*. To locate this document, go to the HP-UX Security docs page at: www.hp.com/go/hpux-security-docs. On this page, select **HP-UX LDAP-UX Integration Software**.

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.

Related Information

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

[HP-UX 11i v3 Networking Software](#).

- *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*

2 ONCplus B.11.31.17

This chapter contains the most recent product information pertaining to the Open Networking Computing (ONC) product, version B.11.31.17, which is supported on the HP-UX 11i v3 operating system. This chapter addresses the following topics:

- “Announcement” (page 6)
- “What’s New in ONCplus B.11.31.17” (page 6)
- “Enhancements and Fixes in ONCplus to date” (page 6)
- “Installation Information ” (page 14)

Announcement

This version of ONCplus (B.11.31.17) is supported on systems running the HP-UX 11i v3 operating system.

What's New in ONCplus B.11.31.17

ONCplus B.11.31.17 is an enhancement and defect fix release.

nfs4cbd daemon port number can be configured using NFS4CBD_PORT variable in `/etc/default/nfs` file. For example, to configure nfs4cbd daemon to run on port 5555, modify `NFS4CBD_PORT = 5555` in `/etc/default/nfs` file. For more details, refer *NFS Services Administrator's Guide*.

All features introduced in previous ONCplus versions are supported in ONCplus B.11.31.17.

Enhancements and Fixes in ONCplus to date

This section lists the enhancements and fixes for ONCplus. It also specifies the version number of ONCplus in which the new feature was introduced and the defect was fixed.

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

Table 1 ONCplus Enhancements and Software Defect Fixes

CR ID	Description	ONCplus Version
QXCR1001238441	nfslog_records_flush_to_disk can cause self deadlock.	B.11.31.17
QXCR1001236356	EBUSY errors during normal unmounting of a NOCTO NFS filesystem.	B.11.31.17
QXCR1001245206	nfssrv module fills up /var if nfslogd failed to process buffer file.	B.11.31.17
QXCR1001260938	11.31 client LM might use wrong protocol while communicating with Srv LM.	B.11.31.17
QXCR1001258028	11.31 deadlock between rfs4_do_open() and rfs4_op_release_lockowner().	B.11.31.17
QXCR1001258012	11.31 deadlock between share, mountd, and nfsd.	B.11.31.17
QXCR1001250258	11.31 10 seconds delay after NFS server failover.	B.11.31.17
QXCR1001264650	Mount using proto=tcp6 does not work	B.11.31.17
QXCR1001240541	11.31 mount -o proto=tcp fails if firewall blocks udp-traffic'	B.11.31.17
QXCR1001233209	11.31 possible data-corruption, when mounting multiple lofs-filestems via NFS.	B.11.31.17
QXCR1001216686	11.31 unable to unshare a deleted directory.	B.11.31.17
QXCR1001270749	11.31 a pv4 hang in lookup.	B.11.31.17

Table 1 ONCplus Enhancements and Software Defect Fixes *(continued)*

CR ID	Description	ONCplus Version
QXCR1001275238	11.31 <code>lm_dup_config</code> panic	B.11.31.17
QXCR1001065684	NFS v4 client fails to do <code>mkdir</code> while a anonymous user or group creates a <code>dir</code> .	B.11.31.17
QXCR1001265094	The <code>da</code> command hang during stress test with NFSv4.	B.11.31.17
QXCR1001277775	11.31 NFSv4-client system panic with large UID.	B.11.31.17
QXCR1001280291	11.31 <code>nfs.server</code> script clears <code>sharetab</code> without <code>unshare</code> .	B.11.31.17
QXCR1001232000	Enable <code>-D_REENTRANT</code> compile option for modules compiled with <code>-lpthread</code> .	B.11.31.17
QXCR1001271923	panic in <code>reg_plist_delete</code> .	B.11.31.17
QXCR1001216448	NFS forced <code>umount</code> fails to release filecache memory leading to memory starvation.	B.11.31.17
QXCR1001190916	NFS deadlock hang between writer application thread and <code>async io</code> threads.	B.11.31.16
QXCR1001202690	11.31 <code>autofs</code> fails to mount NFSv4 over a firewall.	B.11.31.16
QXCR1001209448	<code>rpc.pcnfsd</code> forked many defunct processes.	B.11.31.16
QXCR1001201569	11.31 NFSv4 random gid assigned when create file with <code>O_EXCL</code> set.	B.11.31.16
QXCR1001219901	11.31 with large access list, <code>unshare</code> failed to update <code>/etc/dfs/sharetab</code> .	B.11.31.16
QXCR1001219186	<code>/usr/sbin/rpc.lockd[1332]: t_accept(file descriptor 7/transport tcp) TLI error 0.</code>	B.11.31.16
QXCR1001222312	<code>ls -l</code> fails to report entire directory content.	B.11.31.16
QXCR1001203250	NFS Mounting an LOFS share fails with stale FH when shared with <code>-o fsid</code> option.	B.11.31.16
QXCR1001224990	MC/SG TOC: " <code>cp</code> " command over NFSv4 hogging CPU; <code>hponc_avl_find+0x80</code> .	B.11.31.16
QXCR1001228910	11.31 Ignite performance issue with <code>nfs_enable_ufc_threshold</code> enabled.	B.11.31.16
QXCR1001227362	11.31 NFS client may send wrong FH.	B.11.31.16
QXCR1001229125	11.31 NFS server panic in <code>rfs4_op_rename+0xf70->hpnfs4_vop_getattr+0x30</code> .	B.11.31.16
QXCR1001230603	<code>automount</code> fails with <code>Couldn't bind to reserved port with ports available</code> .	B.11.31.16
QXCR1001221685	11.31 hung <code>umount -f</code> on dead NFS server.	B.11.31.16
QXCR1001228212	ONC admin guide changes for 1303 release.	B.11.31.16
QXCR1001230082	<code>umount -f</code> hang waiting for superpage lock.	B.11.31.16
QXCR1001152584	11.31 possible data-corruption, when mounting multiple <code>lofs</code> -filesystems via NFS.	B.11.31.16
QXCR1001233439	Deadlock between reader application <code>nfs_async_io</code> threads due to <code>mi_async_lock</code> .	B.11.31.16
QXCR1001235838	<code>kctune</code> changes held for next boot for <code>nfs_enable_ufc_threshold</code> throws error.	B.11.31.16
QXCR1001234469	Hang or Panic when multiple threads concurrently do <code>failover_remap</code> .	B.11.31.16
QXCR1001230892	11.31 panics with <code>failover mount</code> and <code>async directio</code> enabled.	B.11.31.16
QXCR1001213468	11.31 <code>nfsd</code> may loop.	B.11.31.16
QXCR1001209969	11.31 NFSv4 <code>callback-portnumber</code> should be configured.	B.11.31.16

Table 1 ONCplus Enhancements and Software Defect Fixes *(continued)*

CR ID	Description	ONCplus Version
QXCR1001240490	va_nodeid may not be populated always under NFSv4.	B.11.31.16
QXCR1001055745	11.31 NFS pv4 readdir performance issue.	B.11.31.16
QXCR1001206222	Contention on the rpcsvc__p_req lock severely degrades SPECsfs performance.	B.11.31.16
QXCR1001133850	11.31 share_nfs(1M) with -log option may cause the system to hang.	B.11.31.15
QXCR1001134401	11.31 NFS KLM-server does not release lock after denied NLM_PROC_GRANTED message.	B.11.31.15
QXCR1001146679	11.31 concurrent share(1M) with -O option may cause double filesystem entries.	B.11.31.15
QXCR1001145872	11.31 concurrent share(1M), may cause double filesystem entries in sharetab.	B.11.31.15
QXCR1001146284	11.31 inconsistent unshare(1M) behavior similar to QXCR1001099945.	B.11.31.15
QXCR1001190907	11.31 nfs.coreX-ipv6 may modify /etc/inetd.conf if swlist(1M) fails.	B.11.31.15
QXCR1001172991	11.31 pcnfsd auth fails if server is trusted and password is >8 characters.	B.11.31.15
QXCR1001166016	Panic in debug IPF kernel for writes on page boundaries.	B.11.31.15
QXCR1001168543	11.31 NFS st_blksize is not consistent.	B.11.31.15
QXCR1001169173	NFS v4 creates deleg state & VHOLDS file without nfsdeleg module.	B.11.31.15
QXCR1001166897	11.31 NFSv4-server hang, generating "vx_iget - inode table overflow" messages.	B.11.31.15
QXCR1001167178	11.31 NFSv4 hang in id_alloc() due to id-bits-array be exhausted.	B.11.31.15
QXCR1001148109	nfsd is not killed when nfs.server is stopped.	B.11.31.15
QXCR1001145478	NFS Thread monopolizing CPUs results in system TOC in Serviceguard Environment.	B.11.31.15
QXCR1001162950	11.31 KLM - a cancel/grant race issue.	B.11.31.15
QXCR1001162529	11.31 KLM4 lm_block_lock has memory leak.	B.11.31.15
QXCR1001189599	Auotofs direct map entry misleads by missing mount options in mnttab.	B.11.31.15
QXCR1001190991	11.31 system panic in hp nfs3_rdwrr/nfs3_async_dio_read.	B.11.31.15
QXCR1001079406	11.31 NFS krb5, sometimes a file created by root at host specified in root=<host> option is set to anon as owner	B.11.31.15
QXCR1001203485	unshare command fails since Share command adds EREMOTE directories in sharetab file.	B.11.31.15
QXCR1001113777	NFS client kitrace records missing for async direct I/O read/write requests.	B.11.31.15
QXCR1001167038	11.31 - NIS client not able to authenticate from passwd.adjunct.byname NIS map.	B.11.31.15
QXCR1001145489	11.31 NFS-client async-write-performance must be improved.	B.11.31.15
QXCR1001180485	NFSV4 changes to loopback hang due to low UFC for QXCR1001059162.	B.11.31.15
QXCR1001214438	Share/Unshare command: Relook.	B.11.31.15
QXCR1001182927	The NFS client may hang for up to 60 seconds after 497 days uptime.	B.11.31.14
QXCR1001181384	On a HP-UX 11i v3 system, the NFS write operation may pause while copying large files when nfs_enable_ufc_threshold = 1.	B.11.31.14
QXCR1001184031	A HP-UX 11i v3 system running ONCplus version B.11.31.13 may panic when nfs_enable_ufc_threshold = 1.	B.11.31.14

Table 1 ONCplus Enhancements and Software Defect Fixes *(continued)*

CR ID	Description	ONCplus Version
QXCR1001025914	The 'ls -ls' command returns an incorrect number of used blocks when it is executed on an NFS system.	B.11.31.13
QXCR1001052992	The statd and automountd daemons incorrectly accept requests sent from a non-loopback address.	B.11.31.13
QXCR1001090660	The gssd daemon may run into a loop under certain conditions.	B.11.31.13
QXCR1001092363	Attempting a forced unmount may result in NFS client system to hang when many write requests are in progress on the system.	B.11.31.13
QXCR1001098196	The NFS client may panic when the nfs_acl_dup_res_impl function is executed.	B.11.31.13
QXCR1001099929	Attempting an "autofs -unmount" may slow down the NFS client system.	B.11.31.13
QXCR1001099945	Adding a host to a netgroup may not work.	B.11.31.13
QXCR1001105312	The pcnfsd daemon may hang when multiple Auth requests are sent to the wlogin function.	B.11.31.13
QXCR1001106195	When the NFS v3 and v4 open/write 32bit API is executed without the O_LARGEFILE parameter, the file size can exceed 2 GB.	B.11.31.13
QXCR1001114093	The nfsmapid daemon may dump core when the DNS server is a windows server.	B.11.31.13
QXCR1001115636	On an NFS protocol v4 system, the default ACLs for the set and get directories may fail with the error "acl failed for file "x",Invalid argument".	B.11.31.13
QXCR1001124632	On an NFS client system, when a single find command is used to find and remove a 0 MB file, the system may panic and display the message "panic: b_spinner_cpsema bad?".	B.11.31.13
QXCR1001127328	The default parameter for the clnt_clts_max_endpoints function has been modified to support the reserving of ports for application usage.	B.11.31.13
QXCR1001127336	The locking mechanism for the e_cnt member of the endpnt_type function has been modified to restrict the maximum number of endpoints.	B.11.31.13
QXCR1001131929	The 'mount -o vers=4,proto=tcp...' command may fail when a firewall is present.	B.11.31.13
QXCR1001044409	A high priority thread may get into a loop by waiting for a lock that is held by a lower priority thread, which may result in a non-pingable hang.	B.11.31.13
QXCR1001059162	NFS client may hang in a loopback environment if UFC filecache is set to very low (say 1%).	B.11.31.13
QXCR1001010420	Executing the open() system call thread over NFSv3 mount may result in system hang.	B.11.31.11
QXCR1001059767	On an HP-UX 11i v3 NFS client system, NFSv3 server-crash recovery does not work.	B.11.31.11
QXCR1001052004	The ls -l command may hang on an NFS-mounted filesystem running ONCplus B.11.31.09.	B.11.31.11
QXCR1001080581	Under certain circumstances, the rpc.lockd daemon mistakenly closes the "listen" fd on the T_DISCONNECT event.	B.11.31.11
QXCR1001090034	NFSv4 server sends the same write verifier across reboots, resulting in data corruption.	B.11.31.11
QXCR1001051156	NFS/VM ASSERT panic occurs during debug kernel testing.	B.11.31.11
QXCR1001055571	Support for reverse name lookup in NFS.	B.11.31.11
QXCR1001058072	NFS server panics when a write request is performed on a file on read-only mount.	B.11.31.11

Table 1 ONCplus Enhancements and Software Defect Fixes *(continued)*

CR ID	Description	ONCplus Version
QXCR1001067886	After an NFS package restart, the NFS client may not be able to access the NFS mount due to stale file handle error (ESTALE).	B.11.31.11
QXCR1001055572	NFS access call returns incorrect ETXTBUSY error.	B.11.31.11
QXCR1001029496	mir_svc_policy_notify - putnext Data TLB Fault panic	B.11.31.10
QXCR1001043662	On HP-UX 11i v3 NFS pv3 client, large hpnfs_var_arena exists due to readdir cache.	B.11.31.10
QXCR1001030379	The readdir_cache is not being freed properly in nfsv3 readdirplus path.	B.11.31.10
QXCR1001015637	HP-UX 11i v2 and HP-UX 11i v3 NFS clients do not distribute load on APA with TCP-mounts.	B.11.31.10
QXCR1001028298	AKRPC clean up and quality improvement.	B.11.31.10
QXCR1001023903	KRPC does not handle RPC_SYSTEMERROR properly which results in panic in an NFS client.	B.11.31.10
QXCR1001018258	NFS client may return EIO error on a TCP hard-mounted NFS-filesystem.	B.11.31.10
QXCR1001043859	The <code>rpcinfo -s</code> command sometimes dumps core.	B.11.31.10
QXCR1001039782	The Thundering Herd issue exists on HP-UX 11i v3.	B.11.31.10
QXCR1001007763	HP-UX 11i v3 NFS server panics in <code>nfs4deleg_support()</code> mode.	B.11.31.10
QXCR1001051998	ONCplus B.11.31.09 - <code>ad</code> command over NFS may not be killed.	B.11.31.10
QXCR1001042152	NFS subsystem-hang is introduced by ONCplus versions B.11.31.09 and B.11.31.09.01.	B.11.31.10
QXCR1001029178	Copy over NFS hangs if the client has small file cache and the <code>nfs3_max_thread</code> tunable is set to 0.	B.11.31.10
QXCR1001028325	The <code>share_nfs(1M)</code> manpage does not explain that NFS-serv shared filesystems cannot be unmounted.	B.11.31.10
QXCR1001004707	HP-UX 11i v3 NFS pv4 readdir request size is larger than <code>nfs4_bsize</code> setting.	B.11.31.10
QXCR1000985674	The <code>rpc.mountd</code> daemon denies mount requests due to incorrect netmask matching.	B.11.31.09.02
QXCR1000986043	Race between <code>autofs</code> lookup and <code>readdir</code> results in busy loop, due to which several threads are kept waiting.	B.11.31.09.02
QXCR1000975336	The <code>nfsd</code> daemon calls <code>nfsys</code> in signal handler without handling <code>errno</code> properly.	B.11.31.09.02
QXCR1001012321	File <code>mmap()</code> with <code>MAP_PRIVATE</code> fails over NFS mount.	B.11.31.09.02
QXCR1001026544	NIS server returns inconsistent data during parallel <code>yp_next()</code> client calls.	B.11.31.09.02
QXCR1001026152	<code>mir_timer()</code> fires after the server disconnects, which sends the client <code>rpcmod</code> to hang state.	B.11.31.09.02
QXCR1001013171	Memory leak on HP-UX 11i v3 NFS PV4 server with delegation enabled.	B.11.31.09.02
QXCR1001008786	NFS v4 server panics during <code>RESTOREFH</code> operation.	B.11.31.09.02
QXCR1001013426	HP-UX 11i v3 NFS PV4 server panics due to double free of <code>pv4 vnode</code> .	B.11.31.09.02
QXCR1001014714	Contention on <code>hpnfs_fsidmap_lock</code> causes significant idle time in <code>SPECsfs</code> .	B.11.31.09.02
QXCR1001014716	Contention on <code>cotsdupreq_lock</code> causes significant idle time in <code>SPECsfs</code> .	B.11.31.09.02
QXCR1001013181	Under certain circumstances the <code>rpc.pcnfsd</code> daemon exits without dumping core.	B.11.31.09.01
QXCR1001011530	The <code>/etc/rc.config.d/nfsconf</code> file has a configuration issue.	B.11.31.09

Table 1 ONCplus Enhancements and Software Defect Fixes *(continued)*

CR ID	Description	ONCplus Version
QXCR1000852734	NFS responds with fixed IP address and not the Serviceguard IP address.	B.11.31.09
QXCR1000929670	The <code>mount_nfs(1M)</code> man page on HP-UX 11i v3 does not explain <code>devs nodevs</code> option.	B.11.31.09
QXCR1000944059	HP-UX 11i v3 memory leak on NFS PV4 client.	B.11.31.09
QXCR1000953573	Unable to clear sharelist when shared directory is deleted.	B.11.31.09
QXCR1000959528	HP-UX 11i v3 <code>ypserv</code> core dump issue.	B.11.31.09
QXCR1000961715	The <code>mount_nfs(1M)</code> man page incorrectly lists "nfs3" as a valid FStype.	B.11.31.09
QXCR1000962354	Duplicate "Start NFS IPv6 subsystem" messages on HP-UX bootup.	B.11.31.09
QXCR1000962396	<code>mir_timer</code> event targeting <code>F_STH_CLOSED</code> NFS TCP Stream.	B.11.31.09
QXCR1000962493	During HP-UX 11i v3 bootup, <code>nfs.core-ipv6</code> may return wrong status message.	B.11.31.09
QXCR1000965447	<code>setoncnv</code> tool must be enhanced to include asynchronous direct I/O tunables.	B.11.31.09
QXCR1000965590	NIS 2.3 gets compilation error for <code>yp_prot.h</code> .	B.11.31.09
QXCR1000967201	NFS does not call <code>alloca()</code> appropriately which may result in a system hang.	B.11.31.09
QXCR1000967219	The <code>rpc.pcnfsd</code> daemon may exit without core dump.	B.11.31.09
QXCR1000973889	Setting the "nfs4_max_threads" tunable to 257 or above causes system panic. (STP)	B.11.31.09
QXCR1000974366	Mount NFS file system parallelly when system starts up.	B.11.31.09
QXCR1000979121	<code>rpc.pcnfsd</code> thread safety issue.	B.11.31.09
QXCR1000981226	Non-root user cannot mount NFSv4 via <code>autofs</code> .	B.11.31.09
QXCR1000982019	HP-UX 11i v3 <code>rpc.pcnfsd</code> fd leaks.	B.11.31.09
QXCR1000982421	The <code>rpc.pcnfsd</code> daemon manipulates effective uid in thread in an unsafe manner.	B.11.31.09
QXCR1000982592	The <code>rpc.pcnfsd</code> daemon exits when it is heavily loaded.	B.11.31.09
QXCR1000987966	Possible memory corruption due to buffer overflow.	B.11.31.09
QXCR1000987972	Incorrect usage of the 'pthread_getspecific' API leads to the <code>rpc.pcnfsd</code> daemon exiting without core dump.	B.11.31.09
QXCR1000930420	Increased GETATTR calls when <code>nfs_enable_write_behind</code> is enabled.	B.11.31.09
QXCR1000846123	Asynchronous NFS Direct I/O performance improvement.	B.11.31.08
QXCR1000892130	The NFS mount structure <code>nfs_args</code> is made public.	B.11.31.08
QXCR1000909967	File locking on NFSv4 client fails with EIO.	B.11.31.08
QXCR1000852435	NFSv4 client subsystem may hang when large number of files are operated by multiple client processes and cause memory pressure.	B.11.31.08
QXCR1000916642	Requests with UID and GID set to -2 does not map to anon value.	B.11.31.08
QXCR1000928323	NFS application may hang under certain circumstances.	B.11.31.08
QXCR1000936917	Application writing to an NFS file from the client system might encounter slow write performance when the file is completely locked.	B.11.31.08
QXCR1000919686	NFS client does not work after 11i v3 NFS server failover.	B.11.31.08
QXCR1000891687	NFS write operation may return EWOULDBLOCK error.	B.11.31.08

Table 1 ONCplus Enhancements and Software Defect Fixes (continued)

CR ID	Description	ONCplus Version
QXCR1000933835	The <code>rpcsvc/rquota.h</code> header file is added back to the ISU bundle.	B.11.31.08
QXCR1000932436	The <code>xdr_setpos()</code> functions incorrectly on ONCplus B.11.31.01 and later versions.	B.11.31.08
QXCR1000862332	NFS client does not do <code>readdir</code> .	B.11.31.08
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 <code>nfs4deleg</code> stackable module.	B.11.31.07
QXCR1000833665	Performance improvement in <code>rpc.statd</code> to improve the crash recovery.	B.11.31.07
QXCR1000886408	Under certain circumstances, NFS client with <code>tcp</code> mount fails to connect to the server, even when the server is available on the network	B.11.31.07
QXCR1000890608	Numerous <code>ypserv</code> processes waiting for <code>lockf()</code> 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 <code>tcp</code> and <code>noac</code> options, setting <code>nfs3_bsize</code> 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 <code>nfs2_srv_read_copyavoid</code> or <code>nfs3_srv_read_copyavoid</code> tunable may cause the following problems: - Memory corruption with the message block getting corrupted - Data corruption during READ, when the <code>nfs3_bsize</code> 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 <code>xdr_callmsg()</code> , <code>xdr_callhdr()</code> , and <code>xdr_replymsg()</code> .	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 <code>clnt_clts_dispatch_notify()</code> .	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 <code>ip6.arpa</code> domain to the <code>getnameinfo(3n)</code> command.	B.11.31.05
QXCR1000847469	Under certain circumstances, ONCplus B.11.31.03 and B.11.31.04 may cause panic.	B.11.31.05
QXCR1000590736	<code>rpc.yppupdated</code> , <code>ypxfrd</code> , and <code>keyserf</code> are not killed when the NIS Master Server is enabled and disabled by <code>smh</code> .	B.11.31.04
QXCR1000593073	Documentation issues in <code>/etc/rc.config.d/namesvrs</code> .	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

Table 1 ONCplus Enhancements and Software Defect Fixes *(continued)*

CR ID	Description	ONCplus Version
QXCR1000779532	Remove nfs4_nnode 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
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 to a very low value.	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 exceeds eight characters in length.	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

Table 1 ONCplus Enhancements and Software Defect Fixes *(continued)*

CR ID	Description	ONCplus Version
QXCR1000591875	Added functionality for the NIS 2.3 client.	B.11.31.02
QXCR1000592166	An AutoFS client using NFSv4 cannot automount an NFS server that does not support NFSv4.	B.11.31.02
QXCR1000592981	The automountd daemon hangs when halting a package.	B.11.31.02
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.yppasswdd.	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 the -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.11.31.17. It also includes information on verifying the installation of ONCplus B.11.31.17.

System Requirements

The following requirements are necessary for installing ONCplus B.11.31.17:

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

Installing ONCplus B.11.31.17

- ❗ **IMPORTANT:** Back up your system before installing the product.
-

To install ONCplus B.11.31.17 on your system:

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

```
swlist | grep ONCplus
```

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

```
ONCplus      B.11.31.16      ONC+ 2.3
ONCplus      B.11.31.15      ONC+ 2.3
ONCplus      B.11.31.14      ONC+ 2.3
ONCplus      B.11.31.13      ONC+ 2.3
ONCplus      B.11.31.12      ONC+ 2.3
ONCplus      B.11.31.11      ONC+ 2.3
ONCplus      B.11.31.10      ONC+ 2.3
ONCplus      B.11.31.09.02   ONC+ 2.3
ONCplus      B.11.31.09.01   ONC+ 2.3
ONCplus      B.11.31.09      ONC+ 2.3
ONCplus      B.11.31.08      ONC+ 2.3
ONCplus      B.11.31.07.01   ONC+ 2.3
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.11.31.17.depot` file, move it to the `/tmp` directory:

```
/tmp/ONCplus_B.11.31.17.depot
```

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

```
swlist -d @ /tmp/ONCplus_B.11.31.17.depot
```

If ONCplus is downloaded correctly, the output will include:

```
# # Bundle(s): #          ONCplus
          B.11.31.17      ONC+ 2.3 # # Product(s) not contained in
a Bundle: #   PHCO_38048          1.0          libc cumulative
patch
```

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

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

```
swinstall -x autoreboot=true -s \          /tmp/ONCplus_B.11.31.17.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.11.31.17 Installation

To verify the ONCplus installation:

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

```
swverify ONCplus
```

If ONCplus B.11.31.17 is successfully installed, the following message is displayed:

```
Verification succeeded
```

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

```
swlist | grep ONCplus
```

The output will be similar to the following:

```
ONCplus      B.11.31.17  ONC+ 2.3
```

Reverting to a Earlier Version of ONCplus

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

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.11.31.06.01, B.11.31.07, B.11.31.07.01, B.11.31.08, B.11.31.09, B.11.31.09.01, B.11.31.09.02, B.11.31.10, B.11.31.11, B.11.31.12, B.11.31.13, B.11.31.14, B.11.31.15, B.11.31.16) 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.



WARNING!

1. Problem

There is a possibility for an NFS client to experience a system hang if you revert to one of the below ONCplus versions:

- ONCplus B.11.31.09
- ONCplus B.11.31.09.01
- ONCplus B.11.31.09.02

Cause

The NFS client could hang under below conditions/configurations:

- Severe file cache memory pressure due to low file cache configuration.
For example, `filecache_min = filecache_max = 1%`
- When the `nfs3_max_threads` tunable is set to "0".

Workaround

The fix for this issue is available in ONCplus B.11.31.10, released in September 2010.

Following are the possible workarounds to handle this issue,

NOTE: All these workarounds must be considered to ensure the system hang does not occur. Reboot the system, after applying all the workarounds.

- Set the below values to the NFS internal variables:

```
# echo "async_commit_wait_multiplier ?W 0xF4240" |  
adb -w -o /stand/vmunix /dev/kmem  
  
# echo "nfs_awcount_multiplier ?W 0x2" | adb -w -o  
/stand/vmunix /dev/kmem
```

- Increase the value of the `nfs3_max_threads` tunable to least 8 using the following command:

```
# kctune nfs3_max_threads=default
```

- If you are doing a manual un-mount, ensure that no application is accessing the NFS mount point at the time of manual un-mount.

For example, if `/mnt/nfs` is an NFS mounted directory, use the following command to verify and un-mount the directory.

```
# fuser -c /mnt/nfs
```

- If you are using AutoFS to un-mount an NFS file system, then disable automatic un-mount by setting a high time out value.

For example, set `AUTOMOUNT_TIMEOUT` to 99999999 in the `/etc/default/autofs` file.

- ### 2. 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 you to install ONCplus B.11.31.08 September 2009 or later version to get the fix for QXCR1000898226.

3. On an HP-UX server with ONCplus B.11.31.07.01 installed, NFSv3 does not fully support large files.
4. 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.
5. 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
6. 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.

7. 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, execute 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: Starting with ONCplus B.11.31.12, the NFS client is enhanced to operate in the SRP environment. So, when ONCplus B.11.31.12 is used in the SRP environment, it is recommended not to revert to an earlier ONCplus version. Doing so may cause the NFS functionality to behave unexpectedly in the SRP environment.

NOTE:

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.
-

8. Problem

- The write operation may be interrupted when very large files are copied over NFS.
- System crash may occur under heavy load conditions.

Solution

This problem has been fixed in ONCplus B.11.31.14.

Workaround

For ONCplus B.11.31.13, the workaround is to use the `forcedirectio` option for the NFS mount in case of loopback mount or cross mounts with very less UFC cache (say 1%) configured on NFS client systems.

Problem

There is a possibility for an HP-UX 11i v3 system to experience a system hang under the following conditions:

- An application writes a huge NFS file to a filesystem using NFS loopback mount.
- Two systems NFS-mount file systems from each other and simultaneously write huge amounts of data to each other.

Causes

On an HP-UX 11i v3 system with NFS loopback mount enabled, there is a possibility for an application to consume the entire file cache memory while writing a huge NFS file. This could lead to a deadlock between the NFS client and NFS server modules.

For example, when the NFS client consumes the entire file cache, the server cannot execute any of the WRITE requests sent by the client because of the non-availability of file cache memory. This can result in non-availability of file cache memory for other mounted file systems on the node.

There is a high possibility for this problem to occur if the system is configured with low file cache memory.

For example, when the `filecache_min` and `filecache_max` parameters are configured in the range of 1% to 5% of the total physical memory.

Solution

Starting from ONCplus B.11.31.15 onwards, a fix has been implemented for both NFS v3 and v4 client versions.

To avoid this issue, the tunables must be set to the following values:

- `kctune nfs_enable_write_behind = 1`
- `kctune nfs_enable_ufc_threshold = 1`
- `kctune nfs_ufc_threshold_percentage = 50`

For more information about the new tunables, see the *Managing NFS and KRPC Kernel Configurations in HP-UX 11i v3* whitepaper

NOTE: All NFS filesystems have to be unmounted if there is any change in these tunables or UFC `filecache_max`. Please note in this version a new tunable `'nfs_ufc_threshold_percentage'` is introduced to cater for both v3 and v4 mounts instead of old tunable `nfs3_ufc_threshold_percentage`. From ONCplus B 11.31.15 version, `nfs3_ufc_threshold_percentage` is made dummy tunable. Start using new tunable `'nfs_ufc_threshold_percentage'` instead of `'nfs3_ufc_threshold_percentage'`.

On ONCplus B.11.31.14 version, Fix was implemented only for NFSv3 client version only.

To enable this fix on this version, following tunables must be set to the below values:

```
kctune nfs_enable_write_behind = 1
```

- kctune nfs_enable_ufc_threshold = 1
- kctune nfs3_ufc_threshold_percentage = 50

NOTE: All NFS v3 filesystems have to be unmounted if there is any change in these tunables or UFC filecache_max.

Workaround

For ONCplus B.11.31.13 and earlier versions, the workaround is to use the "forcedirectio" option for the NFS mount.

Problem

When ONCplus B.11.31.12 is operating in the SRP environment, if the `srp -stop` command is executed while there are active NFS mount points present in the SRP environment, it may hang and a message similar to the following will be logged in the `/var/adm/syslog/syslog.log` file:

```
vmunix: WARNING: SRP 13 shutdown called with 1 active NFS mounts.  
Waiting for unmount to complete.
```

Cause

When the command is executed to shut down the SRP daemon, it may initiate the NFS client shutdown while there are active NFS mount points present in the SRP environment. Due to this, the `srp -stop` command may hang while waiting for the active NFS filesystems to unmount.

Workaround

- Before issuing the command, it is recommended to verify whether the NFS mount point in the SRP environment is active, by executing the `fuser -c` command on the NFS mount point from the global view of SRP.

```
# fuser -c /var/hpsrp/srp01/opt/nfstest
```

where:

```
/var/hpsrp/srp01/opt/nfstest is an NFS mount point under the SRP environment  
srp01
```

This command displays the process ids (PIDs) of the active processes on the NFS mount point.

```
/var/hpsrp/srp01/opt/nfstest: 11352c
```

where:

11352 is the PID of the active process on the NFS mount point

If there are no active processes on the NFS mount point, then execute the command. Otherwise, wait until there are no active processes.

- If the `srp -stop` command has already been executed and it hangs, force unmount all the NFS filesystems under the SRP instance which is being shutdown.

```
# umount -f /var/hpsrp/srp01/opt/nfstest
```

This unblocks the `srp -stop` command and allows it to complete the operation.

NOTE: Doing a force unmount may cause data loss for open files. After the unmount is complete, if a program tries to access files from an unmounted filesystem, it will get an EIO error.

For information on HP-UX SRP commands and HP-UX Containers, see the *HP-UX Containers (SRP) Administrator's Guide*.

To locate this document, go to the following location on the HP Business Support Center:

www.hp.com/go/virtualization-manuals.

On this page, select **HP-UX Containers (SRP) Software**.

3 Features Introduced in previous versions of ONCplus

This chapter describes the features introduced and defect fixes in all the previous versions of ONCplus B.11.31.

- [“Features Introduced in ONCplus B.11.31.16”](#) (page 24)
- [“Features Introduced in ONCplus B.11.31.15”](#) (page 24)
- [“Features Introduced in ONCplus B.11.31.14”](#) (page 24)
- [“Features Introduced in ONCplus B.11.31.13”](#) (page 25)
- [“Features Introduced in ONCplus B.11.31.12”](#) (page 25)
- [“Features Introduced in ONCplus B.11.31.11”](#) (page 25)
- [“Features Introduced in ONCplus B.11.31.10”](#) (page 25)
- [“Features Introduced in ONCplus B.11.31.09.02”](#) (page 26)
- [“Features Introduced in ONCplus B.11.31.09.01”](#) (page 26)
- [“Features Introduced in ONCplus B.11.31.09”](#) (page 26)
- [“Features Introduced in ONCplus B.11.31.08”](#) (page 26)
- [“Features Introduced in ONCplus B.11.31.07.01”](#) (page 27)
- [“Features Introduced in ONCplus B.11.31.07”](#) (page 27)
- [“Features Introduced in ONCplus B.11.31.06”](#) (page 27)
- [“Features Introduced in ONCplus B.11.31.05”](#) (page 28)
- [“Features Introduced in ONCplus B.11.31.04”](#) (page 28)
- [“Features Introduced in ONCplus B.11.31.03”](#) (page 29)
- [“Features Introduced in ONCplus B.11.31.02”](#) (page 29)
- [“Features Introduced in ONCplus B.11.31.01”](#) (page 30)
- [“Features Introduced in ONCplus B.11.31_LR ”](#) (page 30)
- [“Software Availability in Native Languages”](#) (page 35)

Features Introduced in ONCplus B.11.31.16

ONCplus B.11.31.16 is a defect fix release and does not include any new features. All features introduced in previous ONCplus versions are supported in ONCplus B.11.31.16

For information on the defects fixed, see [“Enhancements and Fixes in ONCplus to date”](#) (page 6).

Features Introduced in ONCplus B.11.31.15

ONCplus B.11.31.15 is a defect fix release and does not include any new features. All features introduced in previous ONCplus versions are supported in ONCplus B.11.31.15

For information on the defects fixed, see [“Enhancements and Fixes in ONCplus to date”](#) (page 6).

Features Introduced in ONCplus B.11.31.14

ONCplus B.11.31.14 is a defect fix release and does not include any new features. All features introduced in previous ONCplus versions are supported in ONCplus B.11.31.14

For information on the defects fixed, see [“Enhancements and Fixes in ONCplus to date”](#) (page 6).

Features Introduced in ONCplus B.11.31.13

ONCplus B.11.31.13 is a defect fix release and does not include any new features. All features introduced in previous ONCplus versions are supported in ONCplus B.11.31.13.

For information on the defects fixed, see “Enhancements and Fixes in ONCplus to date” (page 6).

Features Introduced in ONCplus B.11.31.12

ONCplus B.11.31.12 includes the following enhancement. There are no defect fixes included in this release.

Starting with this version, ONCplus has been enhanced to provide support for HP-UX Secure Resource Partitions (SRP) A.03.00 software. All versions of the NFS client protocol namely, NFSv2, NFSv3 and NFSv4 can now operate in the SRP environment.

NOTE: HP recommends you to install the HP-UX Kernel Support for ONCplus patch PHKL_42002, which provides support for the NIS functionality to work with SRP. You can obtain this patch from <http://itrc.hp.com>.

When ONCplus is operating in the SRP environment, the following features/services are supported:

- NFSv2, NFSv3 and NFSv4 client services.
- AutoFS service.
- RPC port mapper service daemon (`rpcbind`): All the RPC user space daemons and libraries can operate in the SRP environment.
- Network Lock manager (NLM) and Network Status Monitor (`rpc.lockd` and `rpc.statd`) client-side services.
- Generating NFS client mount point statistics using the `nfsstat` command.
- NIS server and client services.

When ONCplus is operating in the SRP environment, the following features/services are not supported:

- NFS server services and the `rpc.pcnfsd` daemon
- CacheFS service
- NFSv4 server's delegation stackable FS module

For information on HP-UX SRP A.03.00 software, see the *HP-UX Containers (SRP) Administrator's Guide*.

To locate this document, go to the following location on the HP Business Support Center:

www.hp.com/go/virtualization-manuals

On this page, select **HP-UX Containers (SRP) Software**.

This version of ONCplus supports all features included in previous ONCplus versions.

For information on the defects fixed, see “Enhancements and Fixes in ONCplus to date” (page 6).

Features Introduced in ONCplus B.11.31.11

For information on the defects fixed, see “Enhancements and Fixes in ONCplus to date” (page 6).

Features Introduced in ONCplus B.11.31.10

ONCplus B.11.31.10 is a defect fix release and does not include any new features. All features introduced in previous ONCplus versions are supported in ONCplus B.11.31.10.

From this version onwards, the minimum configurable value for the following NFS kernel tunables has been increased from “0” to “2” for system stability reasons.

- nfs2_max_threads
- nfs3_max_threads
- nfs4_max_threads

NOTE: If the value for any of the above tunables is less than “2”, then the tunable value will be increased to “2” after installation.

For information on the defects fixed, see [“Enhancements and Fixes in ONCplus to date”](#) (page 6).

Features Introduced in ONCplus B.11.31.09.02

ONCplus B.11.31.09.02 is a defect fix release and does not include any new features. All features introduced in previous ONCplus versions are supported in ONCplus B.11.31.09.02.

For information on the defects fixed, see [“Enhancements and Fixes in ONCplus to date”](#) (page 6).

Features Introduced in ONCplus B.11.31.09.01

ONCplus B.11.31.09.01 is a defect fix release and does not include any new features. All features introduced in previous ONCplus versions are supported in ONCplus B.11.31.09.01.

For information on the defects fixed, see [“Enhancements and Fixes in ONCplus to date”](#) (page 6).

Features Introduced in ONCplus B.11.31.09

ONCplus B.11.31.09 includes both defect fixes and new features. All features introduced in previous ONCplus versions are supported in ONCplus B.11.31.09. The following feature was introduced in ONCplus B.11.31.09:

New NFS Feature in ONCplus B.11.31.09

NFS introduces the following feature:

Parallel mount feature for performance improvement.

Prior to ONCplus B.11.31.09, when the system booted up, NFS mounted all the file systems listed in the `/etc/fstab` file serially resulting in considerable boot time. To reduce this time, the parallel mount feature is introduced in ONCplus B.11.31.09. This feature allows NFS to mount multiple file systems parallelly thereby reducing the boot time.

The parallel mount feature has a dependency on the file system changes. The file system changes are being delivered through a EP/NCF release. Following are the EP/NCF patch details:

1. Enhancement Patch: PHCO_40590
Release Date: March 2010
2. NCF Bundle: MountallEnh
Release Date: April 2010

These filesystem patches must be installed for NFS to use the parallel mount feature.

For more information about this feature, see the patch documentation for patch PHCO_40590 and the web page for the NCF Bundle.

Features Introduced in ONCplus B.11.31.08

ONCplus B.11.31.08 includes both defect fixes and new features. All features introduced in previous ONCplus versions are supported in ONCplus B.11.31.08. The following features were introduced in ONCplus B.11.31.08:

New NFS Features in ONCplus B.11.31.08

NFS introduces the following features:

- The NFS mount structure `nfs_args` is made public.

The NFS structure `nfs_args` is made public and available for user application development under `/usr/include/nfs/mount.h` header file. The user applications can make use of this structure to mount NFS shares using `mount()` system call. For more information on the usage of the structure, see the `mount_nfs(1m)` manpage.

- Asynchronous NFS Direct I/O performance improvement.

NFS Direct I/O: NFS Direct I/O is I/O to the NFS mounted files which avoids the file system buffer cache of the operating system. It saves memory and improves performance of the database applications that cache their own data independently. It can be enabled for an NFS mount point, using the `forcedirectio` NFS mount option.

This use of `forcedirectio` option results in synchronous NFS I/O requests to the NFS server. The synchronous I/O operation resulted in a low network utilization on fast networks (10 GB Ethernet). Due to this drawback, a new **Asynchronous NFS Direct I/O** feature is introduced. This option allows the NFS client to send multiple I/O requests to the server without a synchronous wait for individual request. This helps in improving the performance of the database with large sequential I/O.

This feature can be enabled using the following three tunables.

- `nfs3_enable_async_directio_read`
- `nfs3_enable_async_directio_write`
- `nfs3_max_async_directio_requests`

For more information on these tunables, see the *Managing NFS and KRPC Kernel Configurations in HP-UX 11i v3* whitepaper. To locate this document, go to the HP-UX Networking docs page at: <http://www.hp.com/go/hpux-networking-docs>. On this page, select **HP-UX 11i v3 Networking Software**.

NOTE: This feature is disabled by default.

Features Introduced 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.

For information on the defects fixed, see “[Enhancements and Fixes in ONCplus to date](#)” (page 6).

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 “[Enhancements and Fixes in ONCplus to date](#)” (page 6).

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 “[Enhancements and Fixes in ONCplus to date](#)” (page 6).

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 *Group membership expansion: guidelines for deployment* white paper. To locate this document, go to the HP-UX Core docs page at: www.hp.com/go/hpux-core-docs. On this page, select **HP-UX 11i v3**.

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, see the *Introducing Network File systems Version 4 on HP-UX 11i v3* white paper. To locate this document, go to the HP-UX Networking docs page at: <http://www.hp.com/go/hpux-networking-docs>. On this page, select **HP-UX 11i v3 Networking Software**.

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 `cachefswssize` 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 [“Enhancements and Fixes in ONCplus to date” \(page 6\)](#).

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.