Release Notes
ONC+ Networking Enhancements for HP-UX 10.20 (December 1999)

HP Product Number: B6378AA
Supported Hardware: HP 9000 Systems
Part Number/Date: B3782-90875 / December 1999

Table of Contents:

Chapter 1: Introduction

This ONC+ Networking Enhancement software for HP-UX 10.20 provides selected ONC+ capabilities (NFS Version 3.0 Protocol, AutoFS and CacheFS) to the 10.20 HP-UX software base. NFS Version 3.0 Protocol capability is already present in HP-UX 10.30 and later releases.

Objectives of this Release Note

- Inform current users of features in advance of the upcoming major OS release.
- Present alternative methods for accommodating the interface changes (code changes, changes in build process, etc.).
- Document the extent of changes to HP-UX 10.20 that relate to this software update.

This software and documentation are also available from the World Wide Web at URL:

```
http://www.software.hp.com
```

Chapter 2: NFS Version 3 Protocol

Features

NFS Version 3 Protocol supports improved NFS performance by adding the capability to do safe asynchronous writes over NFS mounts. The NFS client keeps track of outstanding uncommitted data to ensure data integrity on the server. NFS Version 3 Protocol also condenses some of the most commonly-used NFS requests into request pairs that travel as a single packet to the server.

NOTE: This ONC+ Networking Enhancement release does not support NFS over TCP. NFS/TCP support will be supported on a future HP-UX 11.X release.

Summary of Change

The only operational change to the NFS subsystem with the addition of the Version 3.0 Protocol is a new switch on the mount command line to force the version of the protocol used to either Version 2 or Version 3. The -o command line option now takes a parameter "vers", which can be set to 2 or 3 depending on which protocol is desired; for example, using the "-o vers=2" option will force the NFS subsystem to mount the requested file system using the Version 2 protocol. After installing this ONC+ Networking software, NFS defaults to using the Version 3 protocol, but automatically falls back to the Version 2 protocol if the remote server does not respond correctly to a Version 3 mount request.

The NFS subsystem now supports a compatibility switch to select the behavior of certain file system APIs that are aware of file system types. See the "Compatibility Switch" section below for details.

Impact

The NFS subsystem may return values which confuse application programs that expect NFS file systems to return a file system type of 1 (or MOUNT_NFS) if the compatibility switch is not enabled. On the other hand, application programs that are aware of the capabilities of NFS Version 3 or CacheFS file systems can use the new return values available in non-compatible mode to make decisions about how to operate.

The ONC+ Networking Enhancements do not include updated header files for the new functionality at this time, due to potential compatibility problems. Once a solution is found to address this problem, updated header files will be made available, most likely in the form of a patch.

Performance

NFS write performance using NFS Version 3 Protocol is substantially better than write performance using the old NFS Version 2 Protocol (as much as 15x in certain cases). However, there are some corner cases in which there will be little or no improvement. The most important one to mention is that under certain circumstances, application programs doing I/O using memory-mapped files may experience performance over NFS Version 3 that is no better than the performance available over NFS Version 2. This is because the NFS subsystem will revert to using synchronous writes in some instances when

memory-mapped I/O is being used in order to avoid deadlocking problems.

Compatibility Switch

The NFS subsystem now has a system-wide compatibility switch to control the behavior of certain file system APIs. The behavior of the stat(), statfs(), and statvfs() functions is affected (and, by virtue of the fact that they use the affected APIs, the fstat(), fstatfs(), and fstatvfs() functions as well).

If the switch is in compatibility mode (that is, compatible with the original 10.20 system behavior - default), return values from the stat(), statfs(), statvfs(), fstat(), fstatfs(), and fstatvfs() system calls are unaffected.

With the switch in non-compatible mode, these calls return different values in the st_fstype field of the "stat" structure returned by (f)stat(), or in the f_fsid field of the "statfs" or "statvfs" structures returned by (f)statfs() and (f)statvfs(), respectively.

The values returned are appropriate to the type of file system being queried. Calls to sysfs() with these values will return "nfs" for NFS Version 2 file systems, "nfs3" for NFS Version 3 file systems, "autofs" for unmounted file systems being monitored by AutoFS, and "cachefs" for CacheFS mounts

NOTE: CacheFS file systems normally return the value of the underlying mount, except for the f_basetype value in the statvfs structure, which will contain the value "cachefs" for the CacheFS file system.

The compatibility switch will only be available on this 10.20 ONC+ Networking Enhancement NFS upgrade; HP-UX 10.30 and 11.0 already implement the non-compatible behavior.

The setting of the compatibility mode switch is controlled by the "onccompat(1M)" command. See the man page accompanying onccompat for details. "onccompat -n" enables the new behavior; "onccompat -c" enables the compatibility behavior. "onccompat -v" displays the current setting without changing its value. The onccompat command must be run by a user with root privilege, since it controls global system settings.

Chapter 3: CacheFS

Features

The CacheFS file system caches data read from another file system (known as the "back" file system) on another, local file system (known as the "front" file system). In the 10.20 ONC+ Networking Enhancement NFS upgrade release, CacheFS only supports NFS as the back file system (the Sun implementation also supports cdfs). The front file system must be HFS.

CacheFS does not cache writes; data written to a cached file is written directly back to the back file system. Writes to a cached file will invalidate the entire file in the cache.

Performance

CacheFS can supply a performance boost to systems that spend a great deal of time accessing stable, read-only data over NFS mounts. For example, application binaries read via a CacheFS mount over NFS will be cached on the local system, and subsequent accesses will be satisfied from the local disk instead of creating a network request.

In the best case, CacheFS can deliver performance approximately 4x better than accessing the same data over NFS. However, actual performance is affected by many factors. For example, with a very fast NFS server, a fast network interface, and a slow, heavily loaded local disk system, CacheFS access may actually be slower than accessing the data over the network.

Compatibility

See the "Compatibility" section under NFS Version 3 Protocol for details on how CacheFS affects system API return values.

CacheFS entries do show up in the mount table, and thus the /etc/mnttab file will reflect cached mount points. Thus, a cached NFS mount point will be reflected by two entries in /etc/mnttab; one for the NFS mount point to a local mount point, and one for the CacheFS mount point from the local NFS mount point. At the moment, the -m option used by Solaris to hide CacheFS mount points is not implemented in HP-UX.

CacheFS does not maintain tight synchronization with the remote server; attribute or data changes done on the server by one client may take some time to propagate to the cache on other clients. For this reason, it is recommended that CacheFS not be used on mounts where data files are frequently updated.

It is also recommended that mount points that are cached with CacheFS should not be backed up with CacheFS enabled. Doing so is extremely wasteful in terms of system resources, since each file must be read across the network and written to local disk even though it is only being written to a backup tape. To back up NFS mount points that are cached, first unmount the cache, or back up the NFS mount point underlying the cache.

Caching a large number of directories, especially by setting up AutoFS to create cached NFS mounts under the /net directory, can result in large numbers of CacheFS daemon (cachefsd) processes being created. This can cause the system to run out of process table entries, which will prevent applications from starting. The recommended way to set up CacheFS is to explicitly identify those remote mount points that should be cached, and use specific mount commands for those mount points.

Chapter 4: AutoFS/Automount

Features

A daemon that automatically mounts/unmounts NFS file systems.

Summary of Change

The 10.20 ONC+ Networking Enhancements replace the existing Automount feature with AutoFS, which performs the same functions, but has a new, more reliable design. AutoFS also supports NFS Version 3 Protocol, whereas the existing 10.20 version of automounter does not. The automount command still exists, but no longer remains running as a daemon after invocation. Instead, an AutoFS daemon, automountd, is started at system boot time to handle automount requests.

Impact

From an operational standpoint, AutoFS functions in the same manner as the old automounter, and the automount command line options and return values are the same.

Compatibility

Any user-written scripts that expect the automount command to remain running as a daemon will have to be updated to either not expect this behavior, or to check if "automountd" is running instead. AutoFS can no longer be shut down by killing the automount process; instead, shut it down by executing "/sbin/init.d/autofs stop". This will unmount all mounted AutoFS file systems, and then kill the automountd process.

The -n, -M, and -tw options to the automount command are not supported at all in AutoFS. The -m and -tm options are not supported, but their behavior can be configured in different ways, by modifying the nsswitch.conf file to get the -m behavior, and by modifying the automount map entries to specify the timeout for the -tm option. The -tl option is accessed using -t.

Another difference between automounter and AutoFS is that AutoFS no longer uses symbolic links to access the mount points; applications that depend on this behavior explicitly will probably not work correctly.

Alternatives

The existing 10.20 automounter can be re-enabled if desired, by replacing the new automount command with the earlier 10.20 version. However, in this configuration automounter will not mount file systems via the NFS Version 3 Protocol, nor will it mount CacheFS file systems.

Chapter 5: rpc.mountd

Features

mountd is an RPC server that answers file system mount requests. It reads file /etc/xtab (described in exports(4)) to determine which directories are available to which machines. It also provides information on what file systems are mounted by which clients.

Summary of Change

rpc.mountd now supports NFS Version 3 Protocol. It will no longer support the -e or -n options which are defined below (from the mountd man page):

```
-e      Exit after serving each RPC request.  When this
        option is used, the inetd security file
        /usr/adm/inetd.sec can control access to RPC
        services.

-n      Exit only if:

             +  portmap dies (see portmap(1M)),
             +  another rpc.mountd registers with portmap, or
             +  rpc.mountd becomes unregistered with portmap.

        This option is more efficient because a new process
        is not launched for each RPC request.  This option is
        the default.
```

Chapter 6: mount (nfs subcommand)

Features

There is now a subcommand of mount/umount for each file system type. The nfs mount/umount subcommands will mount and umount file systems of type nfs or nfs3. All other file system types have their own mount/umount subcommand.

Summary of Change

Here are the new options to the mount command.

```
 vers=n         try mounting the file systems using the version number
                specified by n.  If that version is not available on
                the remote system mount with whatever version is
                available on the remote system.  Known versions at
                this time are 2 and 3.

largefiles      Allow largefile access on the mount point.

nolargefiles    Do not allow largefile access on the mount point.

-F FStype       For the FStype: use nfs.
```

Compatibility

The new mount is compatible with older versions of mount. The new options are not required and using the old options should work the same as they did on previous releases.

Chapter 7: Compatibility With Other Products

HP GlancePlus/MeasureWare

If you have installed the 10.20 ONC+ Networking Enhancement software on a system with the HP GlancePlus or the MeasureWare Agent performance tools installed, you must ensure the version of these products is current or they will not function.

To check Glance or MeasureWare versions, use the command:

```
/opt/perf/bin/perfstat -v
```

or look at the files in /opt/perf/ReleaseNotes directory.

The HP MeasureWare Agent version must be C.01.00 or later, and the HP GlancePlus/UX version must be B.11.01 or later. These product releases include the /opt/perf/bin/midaemon program version B.10.20.15. Earlier versions of the midaemon program will abort when used with 10.20 ONC+ Networking Enhancements because of kernel changes.

SoftBench

The current version of HP SoftBench does not work correctly with CacheFS mounts. Do not cache directories being used by SoftBench.

Chapter 8: Known Problems

Some of these problems may not apply to your specific hardware implementation. Please refer to the Release Note that applies to the product you are installing:

- Workstation Additional Core Enhancement Release Notes (April 1998)
- ONC+ Networking Enhancement Release Notes (April 1998)

NOTE: HP's Systems Administration Manager (SAM) has not been updated to configure the ONC+ features in this ONC+ Networking Enhancements software.

For updated information regarding known problems, see your HP sales representative.

The following is the list of known problems only for the ONC+ Networking Enhancements software:

Symptom:
After cold install from CD, dhcpdb2conf gets error.

Cause:
After cold installing, the system reboots for the final time and runs "set_parms" to setup networking. The system received its IP through a DHCP server. Right after booting and right before "set_parms" is run, the following message is displayed:

```
/sbin/auto_parms, checking network for DHCP server (see
/etc/auto_parms.log)

dhcpdb2conf(ERR):  Couldn't access /etc/resolv.conf for writing
```

After this, "set_parms" runs OK and the system is normal.

Action:
This is an extraneous error which can be ignored.

Legal Notices

Corporate Offices:

Hewlett-Packard Co.
3000 Hanover St.
Palo Alto, CA 94304

This guide’s printing date and part number indicate its current edition. The printing date changes when a new edition is printed. (Minor corrections and updates which are incorporated at reprint do not cause the date to change.) The part number changes when extensive technical changes are incorporated. New editions of this manual will incorporate all material updated since the previous edition.