

HP DECnet-Plus for OpenVMS

Release Notes

January 2005

This manual describes installation notes for this release, new features not discussed in the current documentation, corrections included in this release, documentation errata, and operational notes.

Revision/Update Information:	This is a new manual.
Operating System:	OpenVMS I64 Version 8.2 OpenVMS Alpha Version 8.2
Software Version:	HP DECnet-Plus for OpenVMS Version 8.2

Hewlett-Packard Company
Palo Alto, California

© Copyright 2005 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required 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 HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group.

Printed in the US

Contents

Preface	vii
1 Installation	
1.1 Features Not Supported in This Release	1-1
1.2 License Requirements	1-1
1.2.1 License Requirements for OpenVMS I64 Systems	1-1
1.2.2 License Requirements for OpenVMS Alpha Systems	1-1
1.3 OpenVMS Software Version Requirement	1-2
1.4 Other Software Version Requirements	1-2
1.5 Installing the New Release	1-2
1.6 AUTOGEN Errors During Installation (OpenVMS Alpha Only)	1-3
2 New Features	
2.1 Network Management	2-1
2.1.1 NET\$CONFIGURE Improvements	2-1
2.1.2 NCL Updates	2-1
2.1.3 NCL HELP Updates	2-1
2.1.3.1 Added New Material about Backtranslations and Towers	2-1
2.1.3.2 Added Information about Support for IP Towers	2-2
2.1.3.3 Added Help for register and modify decnet_register Commands	2-2
2.1.3.4 Added Information about Flushing the CDI Cache	2-2
2.1.3.5 Modified Description of the Session Control Naming Search Path Characteristic	2-2
2.1.3.6 Directory System Agent (DSA) Entity Help Updated	2-2
2.1.3.7 Added DNS/BIND Namespace Examples for LOOP Help	2-2
2.1.3.8 Updated Help for FDDI Station Link Event Reporting	2-2
2.2 New Devices Supported by MOP	2-3
2.3 Transport Improvements	2-3
2.4 Host-Based Routing Improvements	2-3
2.5 Session Control Improvements	2-4
2.6 Distributed Naming Service (DECdns)	2-4
2.6.1 DECdns Now Supports Dynamic Device Recognition	2-4
2.6.2 DECdns Now Supports DECnet over TCP/IP Links	2-4
2.6.2.1 Protocol Towers	2-5
2.6.3 decnet_register Now Supports IP Towers in the DECdns Namespace	2-5
2.7 Distributed Time Service (DECdts)	2-6
2.7.1 DECdts Now Supports Dynamic Device Recognition	2-7
2.7.2 DECdts Now Supports DECnet over TCP/IP Links	2-7

3 Corrections Included in This Release

3.1	Network Management	3-1
3.1.1	NET\$MGMT	3-1
3.1.2	NCL\$GLOBALSECTION.DAT	3-1
3.1.3	Common Trace Facility (CTF)	3-1
3.2	Phase V Common Services	3-1
3.3	Network Layer	3-2
3.3.1	End System Routing	3-2
3.3.2	Host Based Routing	3-3
3.4	Transport Layer	3-3
3.4.1	Network Services Protocol (NSP)	3-3
3.4.2	OSI Transport Protocol	3-4
3.4.3	DECnet over TCP/IP (RFC 1859) and OSI over TCP/IP (RFC 1006)	3-4
3.5	Session Control	3-5
3.6	QIO Interface	3-6
3.7	Common Directory Interface (CDI)	3-6
3.8	Distributed Naming Service (DECdns)	3-6
3.8.1	Clerk Changes	3-7
3.8.2	Server Changes	3-7
3.9	Distributed Time Service (DECdts)	3-7

4 Documentation Errata and Addenda

4.1	Current Product Documentation	4-1
4.2	Documentation Update for Version 8.2	4-2
4.3	Network Management and NCL	4-4
4.3.1	Devices and OpenVMS Device Names	4-4
4.3.1.1	CSMA-CD Devices	4-5
4.3.1.2	FDDI Devices	4-6
4.3.1.3	Synchronous Devices	4-7
4.3.1.4	Asynchronous Devices (OpenVMS VAX Only)	4-8
4.3.2	DTR/DTS	4-8
4.4	OSAK Programming	4-8
4.4.1	called_aei Parameter on A-ASSOCIATE and S-CONNECT Indication Events	4-8
4.4.2	ROSE Minimum Workspace Increased	4-9
4.4.3	Using OSAK with OSI over TCP/IP	4-9
4.4.4	New Status Code for osak_send_more and spi_send_more Routines	4-9
4.4.5	New OSAK API Routine: osak_tc_no_sc()	4-9
4.4.6	Handling of Default Transport Templates	4-10
4.4.7	Omitted Parameter Block Parameters	4-10

5 Operational Notes

5.1	Network Management	5-1
5.1.1	CTF Utility Now Requires Process Privileges	5-1
5.1.2	Tuning to Support the Larger CML BYTLM Value	5-1
5.1.3	Upgrading to DECnet-Plus with Systems Running DEC X.25 Client Software (Alpha Only)	5-2
5.1.4	COLLECT Command and Addressing in decnet_migrate Commands	5-2
5.1.5	Decimal Syntax DSPs Not Supported	5-2
5.2	Devices and Data Links	5-3

5.2.1	Receive Buffers Attribute on HDLC LINK Entity Not Implemented (I64 and Alpha Only)	5-3
5.2.2	FDDI Large Packet Support	5-3
5.3	Transport Layer	5-3
5.3.1	NSP Characteristics for X.25 PVCs	5-3
5.3.2	NSP Flow Control Policy and DECnis and x25Gateway Routers	5-3
5.3.3	OSI Transport	5-4
5.4	Session Control	5-4
5.4.1	Transport Precedence in Predominantly Phase IV Networks	5-4
5.4.2	Session Control Back Translation Soft Links	5-4
5.4.3	Restriction When Using Applications Supported Using Cluster Aliases	5-5
5.5	Distributed Naming Service (DECdns)	5-5
5.5.1	Creating Clearinghouses When IP Address is Zero Causes DNS Server Hang	5-5
5.5.2	Server Startup Delay	5-6
5.5.3	Delay in Node Name-to-Address Translations When DECdns Is Using DECnet over TCP/IP Connections	5-6
5.5.4	Problem Creating Clearinghouse Using DECnet over TCP/IP Connections	5-6
5.5.5	DNS\$CONFIGURE Option 4 Does Not Display IP Address	5-6
5.5.6	Limitation on Size of Checkpoint File Records	5-7
5.5.7	PAGEDYN Resources: Clerks and Servers	5-7
5.5.8	Error Message Generated by DECdns Clerk During Startup	5-8
5.5.9	Miscellaneous DECdns Operational Notes	5-8
5.6	Distributed Time Service (DECdts)	5-8
5.6.1	Automatic Time Zone Changes on Rebooting Clusters	5-8
5.7	OSAK Programming	5-8
5.7.1	User Data Size Restriction for Tracing	5-8
5.7.2	OSAK Version 1.1 No Longer Supplied	5-8
5.7.3	Migrating Applications from Version 2.0 to Version 3.0 OSAK API Interface	5-9
5.7.4	Misleading OSAK Error Message: SPM Congestion at Connect Time	5-9
5.8	DECnet Programming	5-9
5.8.1	64-Bit Virtual Address Restrictions	5-9
5.8.2	Backtranslation of IP Addresses	5-9
5.8.3	Applications Connection Requirement	5-9

Tables

4-1	Documentation Update/Release Notes Cross-Reference	4-2
4-2	CSMA-CD Devices	4-5
4-3	FDDI Devices	4-6
4-4	Synchronous Devices	4-7
4-5	Asynchronous Devices (OpenVMS VAX Only)	4-8

Preface

Intended Audience

This manual is intended for all HP DECnet-Plus for OpenVMS Version 8.2 software users. Read this manual before you install, upgrade, or use DECnet-Plus for OpenVMS.

Structure of This Book

This manual contains the following chapters:

- Chapter 1 contains installation notes for this release.
- Chapter 2 describes new features not documented in the DECnet-Plus Version 8.2/7.1 documentation set.
- Chapter 3 lists corrections included in this release.
- Chapter 4 contains documentation errata and additions.
- Chapter 5 discusses operational notes, restrictions, and known problems with the current release.

This manual contains release notes introduced in the current release and notes from previous DECnet-Plus versions that still apply to the new release. Each release note (except those in Chapter 3) contains a margin note indicating in which product version the note first appeared.

Note

Because Alpha installations will be upgrades from the Version 7.3-2 release, these Release Notes contain corrections found in DECnet-Plus for OpenVMS Version 8.1 as well as those found in HP DECnet-Plus for OpenVMS Version 8.2.

For a list of additional documents available that support this version of the DECnet-Plus for OpenVMS software, see the *DECnet-Plus for OpenVMS Introduction and User's Guide*.

Reader's Comments

HP welcomes your comments on this manual or any of the DECnet-Plus documents. Please send comments to either of the following addresses:

Internet **openvmsdoc@hp.com**
Mail Hewlett-Packard Company
 OSSG Documentation Group, ZKO3-4/U08
 110 Spit Brook Rd.
 Nashua, NH 03062-2698

How To Order Additional Documentation

For information about how to order additional documentation and for online versions of most DECnet-Plus documentation, visit the following World Wide Web address:

<http://www.hp.com/go/openvms/doc/>

Conventions

The following conventions are used in this book.

Convention	Meaning
special type	Indicates a literal example of system output or user input. In text, indicates command names, keywords, node names, file names, directories, utilities, and tools. You can abbreviate command keywords to the smallest number of characters that OpenVMS, NCL, DECdns, DECdts and the other utilities accept, usually three characters.
UPPERCASE	Indicates keywords that you enter. You can type the characters in uppercase or lowercase. You can abbreviate command keywords to the smallest number of characters that OpenVMS, NCL, DECdns, DECdts, or the other utilities accept, usually three characters. Uppercase also indicates the names of files, directories, utilities, tools, commands, parameters, and procedures.
<i>italic type</i>	Indicates a variable.
bold	Indicates a new term defined in the text or important information.
Return	Indicates that you press the Return key.
Ctrl/c	Indicates that you press the Control key while you press the key noted by <i>c</i> .
[YES]	Brackets indicate that the enclosed item is a default value in an installation prompt.
{ }	In command format descriptions, indicates you must enter at least one listed element.
\$	Represents the OpenVMS DCL system prompt.

This chapter discusses installation information not found in the current documentation set.

1.1 Features Not Supported in This Release

Version 8.2 X.25 synchronous data links are NOT supported by HP DECnet-Plus for OpenVMS Version 8.2 for OpenVMS I64 systems. They ARE supported by HP DECnet-Plus for OpenVMS Version 8.2 for OpenVMS Alpha systems.

The FTAM component is NOT supported by HP DECnet-Plus for OpenVMS Version 8.2. The OSAK and VT components ARE supported.

1.2 License Requirements

1.2.1 License Requirements for OpenVMS I64 Systems

Version 8.2 HP DECnet-Plus for OpenVMS Version 8.2 for I64 systems requires the Foundation Operating Environment (FOE) license. This license includes the basic function license (DVNETEND). If you intend to configure your system as a DNS server, a DTS server, or use host-based routing, you must obtain a separate DVNETEXT license.

Warning

If you install the software without the DVNETEXT license and then attempt to configure your system as a DNS server, the DNS server will fail and the configuration utility (NET\$CONFIGURE) will exit.

1.2.2 License Requirements for OpenVMS Alpha Systems

Version 7.3-2 HP DECnet-Plus for OpenVMS Version 8.2 for Alpha systems requires one of two licenses: the basic function license (DVNETEND) or the extended function license (DVNETEXT). If you intend to configure your system as a DNS server, you must use the DVNETEXT license.

Warning

If you install the software without the DVNETEXT license and then attempt to configure your system as a DNS server, the DNS server will fail and the configuration utility (NET\$CONFIGURE) will exit.

Warning for Installations over DECnet-Plus V7.3-1

When DECdns server support for Alpha systems was added to DECnet-Plus in Version 7.3-1, the DECdns software mistakenly verified the VAX

Installation

1.2 License Requirements

extended function license (DVNETRTG) instead of the Alpha extended function license. This error has been corrected in this release.

If you have a DECnet-Plus V7.3-1 system currently configured as a DECdns server, do not install this software until you obtain and install the DVNETEXT license. If you fail to install the DVNETEXT license, your DECdns server will not start.

1.3 OpenVMS Software Version Requirement

Version 8.2 HP DECnet-Plus for OpenVMS Version 8.2 for OpenVMS I64 requires Version 8.2 of the OpenVMS I64 operating system. HP DECnet-Plus for OpenVMS Version 8.2 for OpenVMS Alpha requires Version 8.2 of the OpenVMS Alpha operating system.

1.4 Other Software Version Requirements

Version 8.2 If your network includes X.25 connectivity, HP DECnet-Plus for OpenVMS Version 8.2 requires Version 2.0 of the X.25 for OpenVMS product. See the *HP DECnet-Plus for OpenVMS Installation and Configuration* manual for important information about when to install this software during the DECnet-Plus installation and configuration process.

Important

Version 2.0 of the X.25 for OpenVMS product does not support synchronous devices on OpenVMS I64 systems at this time. Therefore, synchronous X.25 data links are not supported by HP DECnet-Plus for OpenVMS Version 8.2 on OpenVMS I64 systems.

If your network includes TCP/IP connectivity and you plan to enable OSI over TCP/IP or DECnet over TCP/IP, HP DECnet-Plus for OpenVMS Version 8.2 requires Version 5.5 of the HP TCP/IP Services for OpenVMS product.

1.5 Installing the New Release

Version 8.2 Install this kit using the POLYCENTER Software Installation utility (PCSI) by logging into the SYSTEM account. The following example is for the OpenVMS I64 platform.

Entering the following command at the DCL prompt (for OpenVMS Alpha systems, use DECnet_OSI instead of DECNET_PLUS):

```
$ PRODUCT INSTALL DECNET_PLUS/SOURCE=file-location
```

For the new software to take effect, you must reboot the system. If the system is a member of a VMScluster, you must reboot the entire cluster.

Note

This kit contains a correction to the CDI\$SHR executable (see CDI\$SHR.EXE in the Corrections section). To ensure that this correction takes effect on a previously configured system, use the following command to flush the CDI cache after you install the kit:

```
$ MCR NCL FLUSH SESSION CONTROL NAMING CACHE ENTRY "*"

```

1.6 AUTOGEN Errors During Installation (OpenVMS Alpha Only)

Version
7.3-2

The PCSI installation process now calls AUTOGEN near the end of the installation (with the parameters GETDATA, SETPARAMS, and CHECK_FEEDBACK). If AUTOGEN returns any error indication, the PCSI installation process displays the following informational text:

```
%PCSI-I-PRCOUTPUT, output from subprocess follows ...
%AUTOGEN-W-REPORT, please read AUTOGEN warnings in SYS$SYSTEM:AGEN$REPORT.REPORT

```

The AUTOGEN report may contain messages similar to the following:

```
** INFORMATIONAL ** - Removing CLU$PARAMS definition made by unknown product
   _UNKNOWN_$ADD_GBLPAGES = 55000      !Set by (unknown)

```

This message is related to the fact that previous versions of DECnet-Plus were incorrectly using the NEWPARAMS.DAT/CLU\$PARAMS.DAT method to pass system parameter information to AUTOGEN. In previous versions, DECnet-Plus did not label its data in the NEWPARAMS.DAT file. The preceding message indicates that AUTOGEN has found an unlabeled system parameter in the existing CLU\$PARAMS.DAT file and is removing it. If the system parameter was set by DECnet-Plus, the newly-created NEWPARAMS.DAT file contains the current system parameter requirements and these will be automatically included in CLU\$PARAMS.DAT. Therefore, you can safely ignore these messages.

This chapter describes new functionality provided by HP DECnet-Plus for OpenVMS Version 8.2. This section also includes new functionality implemented in Versions 7.3-2, 7.3-1, 7.3, 7.2-1, 7.2, and 7.1 which was not documented in the Version 8.2/Version 7.1 documentation set.

2.1 Network Management

This section describes new network management features.

2.1.1 NET\$CONFIGURE Improvements

Version 8.2 The following changes have been made to NET\$CONFIGURE.COM:

- The NET\$CONFIGURE.COM procedure now informs you of the existence of the ISIS\$CONFIGURE.COM procedure. You can use this procedure to configure a routing node that uses the link state routing protocol. Previously, there was no support for configuring a routing node using the link state routing protocol. For more information about using the ISIS\$CONFIGURE.COM procedure, see the new appendix about configuring link state routing in the *HP DECnet-Plus for OpenVMS Installation and Configuration* manual.

2.1.2 NCL Updates

This section describes new features found in NCL.

2.1.3 NCL HELP Updates

Since Version 7.1, several changes have been made to NCL HELP. These changes represent new material about new functionality, new material about existing functionality, and corrections to existing material. The sections that follow describe help changes where some of the material discussed in the HELP updates is not present in the current documentation set or in these Release Notes.

2.1.3.1 Added New Material about Backtranslations and Towers

Version 7.3-2 Added significant explanatory material about managing backtranslation softlink and tower maintenance entities to the help text for the following commands:

```
DELETE SESSION_CONTROL BACKTRANSLATION SOFTLINK
SHOW SESSION_CONTROL BACKTRANSLATION SOFTLINK
UPDATE SESSION_CONTROL BACKTRANSLATION SOFTLINK
UPDATE SESSION_CONTROL TOWER MAINTENANCE
```

For more information, see the help text (of particular interest is the help text for the two update commands).

New Features

2.1 Network Management

2.1.3.2 Added Information about Support for IP Towers

Version 7.3-2 Added a description of an IP tower to the NCL_Introduction -> Syntax -> Data_Types -> TowerSet topic.

Added IP examples to the help text for commands used to manage session control backtranslation softlink and session control tower maintenance entities (see also Section 2.1.3.1).

Added an example of registering an IP tower using the decnet_register node registration tool to the Network_Management -> Tools -> Decnet_Register topic.

2.1.3.3 Added Help for register and modify decnet_register Commands

Version 7.3-2 Added new material about the modify and register commands used to set the session control version.

For more information about why these command examples may be useful, see the help text.

2.1.3.4 Added Information about Flushing the CDI Cache

Version 7.3-2 Added material to the help text for the Flush_(OpenVMS) topic. This topic discusses flushing an entry from the Session Control Naming Cache (known as the CDI cache). Also added an example of flushing an IP cache entry.

2.1.3.5 Modified Description of the Session Control Naming Search Path Characteristic

Version 7.3-2 Modified the description of the naming search path characteristic in the help text for the show session control command to reflect relaxed restrictions in both the content of and the order of the naming search path for DECdns servers. For more information about this topic, see Chapter 7 of the *HP DECnet-Plus for OpenVMS Installation and Configuration* manual (in particular, the section entitled "Using the DECdns Namespace as Secondary on a DECdns Server."

2.1.3.6 Directory System Agent (DSA) Entity Help Updated

Version 7.3-2 Replaced the Directory_Module help topic with new help information for the HP OpenVMS Enterprise Directory product (HP's implementation of the ITU-T X.500 Recommendations for OpenVMS).

2.1.3.7 Added DNS/BIND Namespace Examples for LOOP Help

Version 7.3-1 Added DNS/BIND namespace examples to the help for the loop loopback application command. For example, the example section now includes the following DNS/BIND example:

```
loop loopback application name DOMAIN:LAMCHP.SCH.DEC.COM
```

2.1.3.8 Updated Help for FDDI Station Link Event Reporting

Version 7.3 Updated the Event_Messages -> FDDI_Station -> Link topic to reflect restrictions in FDDI Station Link event reporting which are imposed by the OpenVMS device drivers.

The following events are not logged, but their respective counters are incremented when these events occur:

- block check error
- directed beacon received
- duplicate address test failure
- duplicate token detected
- fci strip error
- frame status error
- link buffer unavailable
- pdu length error

receive data overrun
ring beacon initiated
ring purge error
trace initiated
transmit failure
transmit underrun

The "unrecognized individual pdu destination", "unrecognized multicast pdu destination", and "user buffer unavailable" events can be reported, and their respective counters can increase accordingly. However, DECnet blocks "unrecognized multicast pdu destination" by default. Refer to the `SYS$MANAGER:NET$EVENT_LOCAL.TEMPLATE` file if you wish to pass this event.

The "ring initialization initiated" and "ring initialization received" events are now logged by certain OpenVMS device drivers. If an OpenVMS device driver supports these events, they are not logged on a one-to-one basis with increases in their corresponding counters. Instead, only one error report is issued for any number of these events that occur within a sampling period. This is intended to decrease any possibility of flooding OPCOM.

2.2 New Devices Supported by MOP

Version 7.3

Modified `NET$NISCS_LAA.EXE` to recognize new style device names returned by `$GETDVI`. This permits loads over SCSI PAC devices. Prior to this release, the satellite load failed with the following messages displayed at the satellite:

```
%VMScLuster-I-MSCPCONN, Connected to a MSCP server for the system disk,  
node nnnnnn  
%VMScLuster-E-NOT_SERVED, Configuration change, the system disk is no  
longer served by node nnnnnn
```

2.3 Transport Improvements

Version
7.3-1

This release of the `NET$TPCONS.EXE` module contains the following improvement:

- The DECnet-Plus interface has been redesigned to allow independent restarts between the DECnet OSI Transport module and the HP TCP/IP Services for OpenVMS product. Previously, the entire system had to be rebooted if the TCP/IP PATHWORKS Internet Protocol (PWIP) interface was restarted. Now, if the TCP/IP PWIP interface is disabled and then reenabled, DECnet-Plus starts new transport connections. Similarly, if the DECnet OSI Transport module is disabled and then reenabled, the TCP/IP PWIP interface receives new DECnet transport connections through the new TPCONS ports.

2.4 Host-Based Routing Improvements

Version 8.2

This release supports the configuration of host-based routers that use the link state routing algorithm. The `NET$CONFIGURE` procedure always creates routers that use the routing vector algorithm. However, a new command procedure (`ISIS$CONFIGURE.COM`) can be used to modify the router configuration to use the link state routing algorithm. This procedure also allows more control of the routing configuration than is found in `NET$CONFIGURE.COM`. For more information, see the "Host-Based Routing" section in Chapter 8 of the *HP DECnet-Plus for OpenVMS Network Management*

New Features

2.4 Host-Based Routing Improvements

manual and the new ISIS\$CONFIGURE appendix in the *HP DECnet-Plus for OpenVMS Installation and Configuration* manual.

This release of the NET\$ROUTING_IS.EXE module contains the following improvements:

- Version 8.2* • Raised the maximum allowed routing circuits from 16 to 32.
- Version 7.3* • Modified logic used to set maximum circuit segment size for FDDI circuits. Previously, the calculation could sometimes use the Phase IV default of 576. This has been changed to always use the FDDI default of 4352.

2.5 Session Control Improvements

This release of the NET\$SESSION_CONTROL.EXE module contains the following improvement:

- Version 7.3* • On I64 and Alpha systems, the NET\$SESSION_CONTROL.EXE module now uses the name string from the PSB instead of GETJPI.

2.6 Distributed Naming Service (DECdns)

The following sections describe new features of DECdns.

2.6.1 DECdns Now Supports Dynamic Device Recognition

Version 8.2 The static device tables formerly used to determine the devices used by DECdns have been removed. Now, DECdns uses the \$DEVSCAN and \$GETDVI system services to build a list of devices that have the following characteristics:

- a device class of DC\$SCOM (synchronous communication device)
- a device characteristic of DEV\$V_NET
- a device status of UCB\$V_ONLINE and UCB\$V_TEMPLATE
- a device name in the form `_xxy0:`

You can use the logical name DNS\$ETHERNET_DEVICE to provide a list of devices that DECdns should NOT use. All devices must be in the form `_xxy0:.` The string can contain spaces and other text which is ignored by DTSS. For example, the following command tells DECdns not to use the `_EIA0:` and `_FWA0:` devices.

```
$ DEFINE/SYSTEM DNS$ETHERNET_DEVICE "Don't use _EIA0: and _FWA0: "
```

Note

Unterminated network adapters can cause the dynamic device recognition process to hang. Either terminate all network adapters or include any unterminated devices in the DNS\$ETHERNET_DEVICE logical definition.

2.6.2 DECdns Now Supports DECnet over TCP/IP Links

Version 7.3–2 DECdns has been enhanced to support the DECnet over TCP/IP transport. The sections that follow discuss various topics related to this new feature. For more information, see the *HP DECnet-Plus for OpenVMS DECdns Management* manual.

2.6.2.1 Protocol Towers

The `DNA$Towers` attribute of the object entry for a Phase V node can now include a protocol tower with an IP address. The third tower in the following example is an IP protocol tower.

```
DNS>show obj .xpr.wzbang dna$towers
          SHOW
          OBJECT WINTER:.xpr.wzbang
          AT 09-MAR-2004:14:55:02
          DNA$Towers (set) = :
          Tower 1 Floor 1 = 01 13 (null)
          Tower 1 Floor 2 = 03 00 13
          Tower 1 Floor 3 = 04 (null)
          Tower 1 Floor 4 = 06 49 00 18 aa 00 04 00 2a 60 20
          DNA$Towers (set) = :
          Tower 2 Floor 1 = 01 13 (null)
          Tower 2 Floor 2 = 03 00 13
          Tower 2 Floor 3 = 05 de c0
          Tower 2 Floor 4 = 06 49 00 18 aa 00 04 00 2a 60 21
          DNA$Towers (set) = :
          Tower 3 Floor 1 = 01 13 (null)
          Tower 3 Floor 2 = 02 00 13
          Tower 3 Floor 3 = 05 de c0
          Tower 3 Floor 4 = 09 a1 72 5e 45
DNS>
```

2.6.3 decnet_register Now Supports IP Towers in the DECdns Namespace

Version
7.3-2

Normally, the IP towers and IP backtranslation softlinks for all nodes are now automatically updated in the DECdns namespace during regular tower and softlink maintenance operations. Occasionally, it may be necessary to perform manual updates or to verify that the current information is correct. For this reason, the `decnet_register` tool has been enhanced to show and register IP towers and backtranslation softlinks for nodes within the DECdns namespace. The following example illustrates how to use menu Option 1 to display information about the node `WINTER:.XPR.CASE`:

```
$ run decnet_register

DECNET_REGISTER - Manage node registrations in network directory services
Use Return, CTRL/N, and CTRL/P to move between input fields
Use "?" to obtain help, CTRL/Z to cancel

 1 - Show information about registered node names
 2 - Register or modify node names
 3 - Update registered node towers using information from the nodes
 4 - Rename a registered node name
 5 - Repair the synonym and address links for registered node names
 6 - Deregister node names

 7 - Export node names to a data file
 8 - Import node names from a data file

 9 - Set preferences and network values

10 - Manage the directory service
11 - Spawn to DCL

* Option (use CTRL/Z to exit): 1

Show registered node information
Use Return, CTRL/N, and CTRL/P to move between input fields
Use "?" to obtain help, CTRL/Z to cancel
```

New Features

2.6 Distributed Naming Service (DECdns)

Specify the directory service as LocalFile, DECdns, or PhaseIV.

* Directory service: DECdns

Specify the node to show using an explicit or wildcard name, an NSAP, or a Phase IV synonym or address.

* Node name or address: WINTER:.XPR.CASE

Specify the information to display as either brief, full, or names. Specify the output file name (a blank line indicates the terminal).

* Display format: full

* Output file:

Press Return to show the node values, CTRL/Z to cancel

Directory Service: DECdns

Node name: WINTER:.XPR.CASE

Phase IV synonym: CASE

Address tower protocol and selector values:

```
Session: DNA_SessionControlV3 (SC3)
          00 13
Transport: DNA_OSITransportV1 (TP4)
          DE C0
Routing:  DNA_OSInetwork (CLNS)
          49::00-18:AA-00-04-00-24-60:21 (24.36)
```

```
Session: DNA_SessionControlV3 (SC3)
          00 13
Transport: DNA_NSP (NSP)
          (no selector value)
Routing:  DNA_OSInetwork (CLNS)
          49::00-18:AA-00-04-00-24-60:20 (24.36)
```

```
Session: DNA_SessionControlV2 (SC2)
          00 13
Transport: DNA_OSITransportV1 (TP2 or TP0)
          DE C0
Routing:  DNA_IP (IP)
          161.114.94.62
```

Synonym softlink name:
WINTER:.DNA_NodeSynonym.CASE

Reverse address mapping (backtranslation) softlink names:
WINTER:.DNA_BackTranslation.%X49.%X0018.%XAA0004002460
WINTER:.DNA_BackTranslation.%XA1725E3E

Number of nodes reported on: 1

2.7 Distributed Time Service (DECdts)

The following sections describe new features of DECdts.

2.7.1 DECdts Now Supports Dynamic Device Recognition

Version 8.2 The static device tables formerly used to determine the devices used by DECdts have been removed. Now, DECdts uses the \$DEVSCAN and \$GETDVI system services to build a list of devices that have the following characteristics:

- a device class of DC\$SCOM (synchronous communication device)
- a device characteristic of DEV\$V_NET
- a device status of UCB\$V_ONLINE and UCB\$V_TEMPLATE
- a device name in the form _xxy0:

You can use the logical name DTSS\$ETHERNET_DEVICE to provide a list of devices that DECdts should NOT use. All devices must be in the form _xxy0:. The string can contain spaces and other text which is ignored by DTSS. For example, the following command tells DECdts not to use the _EIA0: and _FWA0: devices.

```
$ DEFINE/SYSTEM DTSS$ETHERNET_DEVICE "Don't use _EIA0: and _FWA0: "
```

Note

Unterminated network adapters can cause the dynamic device recognition process to hang. Either terminate all network adapters or include any unterminated devices in the DTSS\$ETHERNET_DEVICE logical definition.

2.7.2 DECdts Now Supports DECnet over TCP/IP Links

Version 7.3-2 DECdts now supports operation in the DECnet over TCP/IP environment without requiring the use of a configuration file. This requires the use of the DECdns version that supports DECnet over TCP/IP. From the management standpoint, global DECdts servers work much the same as they do in the DECnet-only environment. You still use the advertise command to add a global server to the global set. For more information, see the *HP DECnet-Plus for OpenVMS DECdts Management* manual.

Corrections Included in This Release

This chapter discusses corrections included in HP DECnet-Plus for OpenVMS Version 8.2. The chapter contains an abbreviated list of the updates included in this release. If a specific update you are looking for is not listed, this does not mean that it is not included in this release. For further information on specific updates, contact your HP support representative.

3.1 Network Management

This section describes network management corrections in this release.

3.1.1 NET\$MGMT

The following corrections have been made to the NET\$MGMT graphical user interface utility:

- Modified NET\$MGMT's internal tables to support more than 500 ports. Previously, the tables were hard-coded to support 500 ports; they are now allocated dynamically.

3.1.2 NCL\$GLOBALSECTION.DAT

The following change has been made to NCL\$GLOBALSECTION.DAT:

- Added CMIP code 55 for "OpenVMS I64" systems. Changed CMIP code 28 for "OpenVMS AXP" systems to "OpenVMS Alpha" systems. Changed CMIP code 20 from "Compaq DECnet-Plus for OpenVMS" to "HP DECnet-Plus for OpenVMS."

3.1.3 Common Trace Facility (CTF)

The following correction has been made to the CTF\$DCP.EXE image:

- Modified the code to prevent a possible system crash in CTF\$DISPATCH_KERN by correcting deallocation logic for trace records and CTF work queue elements in CTF\$DISPATCH_KERN and CTF\$SRV\$UNDECLARE.

3.2 Phase V Common Services

This section describes changes made to the common services component of DECnet.

The following correction has been made to the NET\$PROCESS_EMAA.EXE image:

- Corrected a system crash seen while managing a MOP entity. A block deallocation routine was running at an elevated IPL (IOLOCK8) when a page fault occurred. The elevated IPL is not required and was removed, thus allowing the page faulting to occur at a lowered IPL.

Corrections Included in This Release

3.2 Phase V Common Services

- Corrected a system crash caused by a problem with buffer allocation and initialization in the NET\$EMAA_Response routine. The buffer consists of a fixed length control block followed by a variable length output string. The fixed length portion includes a pointer to the variable length output string. Previously, the code was not properly initializing the output string pointer when the output string length was zero. The code has been modified to properly set up a null pointer when the output string length is zero.

The following corrections have been made to the SYS\$NETWORK_SERVICES.EXE image:

- Updated the DECNET_VERSION item code to "0051200".
- Corrected a system crash seen when using the NCL SHOW SESSION CONTROL PORT * command. The code was incorrectly releasing an internal data structure (but not the pointer to the structure) during port takedown while processing was in progress for the NCL command. The released data structure was still referenced during the NCL processing. If the deallocated memory was rapidly reused, this led to the system crash.
- Corrected a system crash seen when CTF tracing was active during a shutdown of the OSI transport. During transport shutdown, the CTF trace information buffer is deallocated. The common code for handling trace terminations, executed after the transport shutdown code, failed to verify the existence of the trace information buffer before using a pointer within the buffer. This code has been modified to properly verify the CTF buffer's existence before using any data within the buffer.

3.3 Network Layer

This section describes corrections made to the network layer.

3.3.1 End System Routing

The following correction has been made to the NET\$ROUTING_ES.EXE image:

- Added a check for the supplied MAC address pointer to the LAN circuit receive completion code. Previously, if the MAC address pointer was invalid a system crash occurred. Now, the VCRP containing the invalid MAC address pointer is simply discarded.
- Prevent a SSRVEXCEPT system crash by ensuring that an FDDI circuit is enabled before sending the last chance ES_HELLO (for all best circuits).
- Corrected a ACCVIO system crash seen when disabling alias ports. The code was not handling the case where there were no ports to disable.
- Corrected a problem with the integrity checking of incoming ES Hello messages. The verification routine that checked the length of the message was using the first byte of the message instead of the length byte preceding the message. This has been corrected. In addition, if the Hello message was found to exceed the 20-byte architectural limit the routine caused a system crash. The routine now returns status to its callers indicating that the message has failed verification. The calling routines now handle this status using appropriate error handling instead of causing a system crash.

3.3.2 Host Based Routing

The following corrections have been made to the NET\$ROUTING_IS.EXE image:

- Corrected several problems that prevented OSI-only host based routing (that is, Phase V link state routing) from operating correctly. Previously, this configuration still required remnants of the Phase IV routing vector algorithm while not working correctly if certain Phase IV routing vector parameters were set.
- Corrected a problem seen when an alias port is disabled. After the disable, the NSAPs associated with the alias port are still listed as a valid routing adjacency and data PDUs are still accepted. However, because the NSAPs are no longer available, several minutes of address unreachable pdu discard events (with a discard reason of destination address unknown) were seen. These messages continued until the adjacency's holding timer expired at which time the adjacency was finally deleted. The problem was corrected by modifying the way the routing module responds to the notice it receives from the alias module that the alias has been disabled. The routing module now sends out an ES Hello message with the alias NSAPs marked with a holding time of zero. This causes routers to immediately delete the adjacency for the alias.
- Deleted unnecessary Phase IV parameter checks when in an pure OSI host-based routing environment.
- Modified the output of the NCL SHOW ROUTING ALL command to omit all Phase IV parameters in a pure OSI host-based routing environment. The following portion of the SHOW ROUTING ALL display is now suppressed in a pure OSI host-based routing environment:

```
Segment Buffer Size           = 0
PhaseIV Buffer Size           = 0
PhaseIV Address               = 0.0
PhaseIV Prefix                = <Default value>
PhaseIV Maximum Address      = 0
PhaseIV Maximum Hops         = 0
PhaseIV Maximum Cost         = 0
PhaseIV Broadcast Routing Timer = 0
PhaseIV Maximum Visits       = 0
```

3.4 Transport Layer

This section describes changes made to the transport layer.

3.4.1 Network Services Protocol (NSP)

The following corrections have been made to the NET\$TRANSPORT_NSP.EXE image:

- Corrected several problems that prevented OSI-only host based routing (that is, Phase V link state routing) from operating correctly. Previously, this configuration still required remnants of the Phase IV routing vector algorithm while not working correctly if certain Phase IV routing vector parameters were set.
- Corrected a system crash caused by a fork to 0 seen during retransmissions on high traffic LAN connections. The problem occurred when the LAN driver was completing the initial transmission request while NSP had already timed out and was trying to retransmit the data. The LAN drivers were modified when fastpath support was added. One of the changes was to use a variety

Corrections Included in This Release

3.4 Transport Layer

of spinlocks instead of IOLOCK8 (as NSP does). This change allowed both processes to modify the scratch area (which contains the fork address) of the transmit VCRP at the same time. NSP has been modified so that it does not modify the VCRP scratch area once the VCRP is given to the LAN drivers.

- Corrected a problem resulting in poor throughput when NSP was using piggyback ACKs. The code now sends a standalone ACK if no data segment is queued for transmission.

3.4.2 OSI Transport Protocol

The following corrections have been made to the NET\$TRANSPORT_OSI.EXE image:

- Corrected a system crash caused by a fork to 0 seen during retransmissions on high traffic LAN connections. The problem occurred when the LAN driver was completing the initial transmission request while OSI transport had already timed out and was trying to retransmit the data. The LAN drivers were modified when fastpath support was added. One of the changes was to use a variety of spinlocks instead of IOLOCK8 (as OSI transport does). This change allowed both processes to modify the scratch area (which contains the fork address) of the transmit VCRP at the same time. OSI transport has been modified so that it does not modify the VCRP scratch area once the VCRP is given to the LAN drivers.
- Corrected a session hang problem seen when a disconnect was received from the remote system and was being processed at the same time that a request to deaccess the NET channel came from the local system. Added a check in the ACT\$Closed_T_DISreq routine, to determine if a remote disconnect has been received and if so then complete the request immediately since the connection takedown is already in progress.
- Corrected a INVEXCEPTN system crash by preventing a double deallocation of the NCCB seen during the disabling of the OSI transport when DECnet is shutting down.

3.4.3 DECnet over TCP/IP (RFC 1859) and OSI over TCP/IP (RFC 1006)

The following corrections have been made to the NET\$TPCONS.EXE image:

- Corrected a forced crash dump caused during network shutdown. The Deactivate TPCB routine was incorrectly called before the connection reached the UNBND state. This occurred because the calling code had attempted to UNBIND the PWIP driver at the same time that the PWIP driver was shutting down. This resulted in an ERROR_ACK to the UNBIND request which was not properly handled.
- Corrected a problem in handling an ERROR_ACK response from the PWIP driver during a BIND sequence. The problem was caused by a reference to an uninitialized variable in the ERROR_ACK handler. The reference caused a system crash in the TPCONS\$DeallocateTPCB routine.
- Corrected a problem with the display of OSI over TCP/IP parameters in the SHOW TRANSPORT PORT ALL STATUS command:

Corrections Included in This Release 3.4 Transport Layer

```
NCL> SHOW OSI TRANSPORT PORT port-name ALL STATUS
Node 0 OSI Transport Port OSI$PORT_0_0050
at 2004-10-29-11:01:51.797-05:00I0.173

Status

    UID                               = A56E02B5-F7E2-11D7-AD31-454453434C55
    .
    .
    Local RFC1006 Port Number         = 0
    Local RFC1006 IP Address          = <Default value>
    Remote RFC1006 Port Number        = 0
    Remote RFC1006 IP Address         = <Default value>
```

Note that the Port Number attribute always displayed as 0 and the IP Address attribute always displayed as "<Default value>". The OSI transport module did not have a method to retrieve this information from the OSI over TCP/IP process. This problem has been corrected and the port and address information is now displayed.

- Corrected a system crash seen during processing of an incoming disconnect when the local system is also initiating a disconnect. The code handling the incoming disconnect detected an invalid control block as a result of the locally-initiated disconnect and caused a system crash. The code handling the incoming disconnect has been modified to correctly pass the disconnect to the PATHWORKS Internet Protocol (PWIP) driver instead of cause a system crash.
- Removed sanity checking in DeallocateTPCB that was no longer compatible with newer HP TCP/IP Services for OpenVMS releases.

3.5 Session Control

The following corrections have been made to the NET\$SESSION_CONTROL.EXE image:

- Corrected a rare SMP issue with cleanup of orphan sessions. Occasionally, a connect reject response is in the process of being delivered to the caller before the transport takedown has completed. In this case, the internal transport takedown event is completed directly to the caller instead of the connect reject. Because session control completion code does not expect this type of event, a crash occurs. The problem has been fixed by waiting for the transport takedown to complete before completing the connect reject to the user.
- Modified the towers created for IP addresses in the DECdns namespace to always use SC2. When support was added for storage of IP addresses, the Session Control value was forced to SC3. Also, Session Control now maintains tower and backtranslation softlinks for both SC2 and SC3. Previously, it maintained softlinks for SC3 only. See also Section 3.7 and Section 3.8.2.
- Corrected a problem seen when a user program received notification of I/O completion using LOCAL event flags. Occasionally, the event flag was set before the IOSB and the data buffer were properly set.

Corrections Included in This Release

3.6 QIO Interface

3.6 QIO Interface

The following corrections have been made to the NET\$DRIVER.EXE, NET\$OSDRIVER.EXE, and NET\$OSVCM.EXE images:

- Added spinlock synchronization to two routines responsible for determining that a NET device's connection count has gone to zero. Before the spinlock synchronization, each routine could independently detect the condition and attempt to close the association at the same time. This resulted in a hung NET device.
- Improved the reporting of connection timeouts in the transmit and receive queue processing code. The code now returns the traditional, more specific timeout error status SS\$_LINKEXIT instead of the generic error status SS\$_FILNOTACC.
- Corrected a problem seen when a user program received notification of I/O completion using LOCAL event flags. Occasionally, the event flag was set before the IOSB and the data buffer were properly set.
- Corrected a system crash seen when multiple deassign calls were made for a connection instance. For example, a remote system could request a disconnect at the same time as the local system. The second call incorrectly modified the control block and returned an error to the caller. This caused the connection instance to remain in the RWAST state. The code has been corrected to allow the disconnect to proceed.
- Corrected a problem with the OS device driver used by OSI applications. The problem was seen only with incoming sessions. If the remote session partner and the device driver simultaneously attempted to close the session, both code paths attempted to deallocate the same internal data structures. This led to a possible hang on the deassign operation. The code has been corrected to properly handle these two events occurring simultaneously.

3.7 Common Directory Interface (CDI)

The following corrections have been made to the CDI\$SHR.EXE image:

- Corrected a problem with Domain proxies brought about by using Session Control V3 (SC3) with RFC1006 connections. The problem was corrected by reverting IP address towers back to SC2 (see also Section 3.5 and Section 3.8.2). CDI is now used to look up the node name for RFC1006 connections. This can cause rejected connections if the backtranslation entries in DECdns are incorrect. To ensure that this change takes effect, use the following command to flush the CDI cache after you install this kit:

```
$ MCR NCL FLUSH SESSION CONTROL NAMING CACHE ENTRY "*"
```

3.8 Distributed Naming Service (DECdns)

This section describes corrections made to the components of the Distributed Naming Service (DECdns).

Corrections Included in This Release

3.8 Distributed Naming Service (DECdns)

3.8.1 Clerk Changes

The following correction has been made to the clerk image (DNS\$ADVER.EXE):

- Corrected a problem that caused a subsequent device to remain uninitialized if a lower device of the same type was offline or unavailable.
- Corrected a problem that could cause more than three copies of the DNS\$CHFAIL.LOG to be present on the system. This condition occurred when the file version of DNS\$CHFAIL.LOG cycled through the upper version boundary (that is, cycling from version 32767 back to version 1). The code now correctly handles this situation and maintains three copies of the file. File versions are rolled over when the file version reaches 32000. For example, if versions 31998, 31999, and 32000 are present, the code renumbers the files to version 1, 2, and 3, respectively.
- Corrected a problem that could cause multiple copies of the file used to maintain the current cache file version (SYS\$SYSTEM:DNS\$CACHE.VERSION) to be present on the system. This condition occurred when the file version of DNS\$CACHE.VERSION cycled through the upper version boundary (that is, cycling from version 32767 back to version 1). The code now correctly handles this situation and maintains a single version of the file.

3.8.2 Server Changes

The following corrections have been made to the server image (DNS\$SERVER.EXE):

- Corrected a problem with Domain proxies brought about by using Session Control V3 (SC3) with RFC1006 connections. The problem was corrected by reverting IP address towers back to SC2 (see also Section 3.7 and Section 3.5). CDI is now used to look up the node name for RFC1006 connections. This can cause rejected connections if the backtranslation entries in DECdns are incorrect.

3.9 Distributed Time Service (DECdts)

The following corrections have been made to the DTSS\$SERVICE.EXE image:

- Corrected a problem that caused a subsequent device to remain uninitialized if a lower device of the same type was offline or unavailable.
- Corrected a problem that caused programs using the JAVA runtime code to go into a COM state loop when the TDF change occurs. A TDF change notification event was added. Also, there are some VMS/THREADS corrections that are related to this problem.

Documentation Errata and Addenda

This chapter contains corrections and additions to the current documentation set for HP DECnet-Plus for OpenVMS.

4.1 Current Product Documentation

Version 8.2 The current documentation set for DECnet-Plus is shown in the following list:

- *DECnet-Plus for OpenVMS Introduction and User's Guide*
Version 7.1
(Order number: AA-Q191F-TE)
- *DECnet-Plus Planning Guide*
Version 7.1
(Order number: AA-QBTHE-TE)
- *HP DECnet-Plus for OpenVMS Installation and Configuration*
Version 8.2
(Order number: BA406-90001)
- *HP DECnet-Plus for OpenVMS Network Management*
Version 8.2
(Order number: BA406-90006)
- *HP DECnet-Plus for OpenVMS Network Control Language Reference*
Version 8.2
(Order number: BA406-90004)
- *DECnet-Plus Problem Solving*
Version 7.1
(Order number: AA-Q18ZE-TE)
- *HP DECnet-Plus for OpenVMS DECdns Management*
Version 8.2/DECdns Version 2.0
(Order number: BA406-90002)
- *HP DECnet-Plus for OpenVMS DECdts Management*
Version 8.2/DECdts Version 2.0
(Order number: BA406-90003)
- *DECnet-Plus FTAM and Virtual Terminal Use and Management*
Version 7.1/FTAM Version 3.2
(Order number: AA-Q78TD-TE)
- *DECnet-Plus for OpenVMS Programming*
Version 7.1
(Order number: AA-Q195E-TE)
- *DECnet-Plus DECdts Programming*
Version 7.1/DECdts Version 1.1
(Order number: AA-PHEQC-TE)

Documentation Errata and Addenda

4.1 Current Product Documentation

- *DECnet-Plus OSAK Programming*
Version 7.1/OSAK Version 3.0
(Order number: AA-PJ1CE-TE)
- *DECnet-Plus OSAK Programming Reference*
Version 7.1/OSAK Version 3.0
(Order number: AA-PJ1DD-TE)
- *DECnet-Plus OSAK SPI Programming Reference*
Version 7.1/OSAK Version 3.0
(Order number: AA-QRAYB-TE)
- *DECnet-Plus FTAM Programming*
Version 7.1/FTAM Version 3.2
(Order number: AA-PJ1EF-TE)
- *DECnet/OSI for VMS CTF Use*
Version 5.5
(Order number: AA-PHEBB-TE)

4.2 Documentation Update for Version 8.2

Several manuals have been updated for HP DECnet-Plus for OpenVMS Version 8.2. See the list in Section 4.1 for the version numbers associated with each manual in the documentation set. Much of the updated material came from entries in the Release Notes for several previous versions of DECnet-Plus. Table 4–1 provides a cross-reference between former Release Note topics and their locations in the new manuals. The manual references are given in the form *manual-chapter or appendix* where *manual* is the following:

INS - *HP DECnet-Plus for OpenVMS Installation and Configuration*
MNG - *HP DECnet-Plus for OpenVMS Network Management*
NCL - *HP DECnet-Plus for OpenVMS Network Control Language Reference*
DNS - *HP DECnet-Plus for OpenVMS DECdns Management*
DTS - *HP DECnet-Plus for OpenVMS DECdts Management*

In many cases, the information in the Release Notes was interleaved with existing information. Therefore, references to specific sections in the new manuals is not feasible.

Table 4–1 Documentation Update/Release Notes Cross-Reference

Release Note Topic	Refer to ...
New Features	
NET\$CONFIGURE	
New option to configure DNS servers	INS-7
DNS server node can have other name service as primary	INS-7
Phase IV migration improvements	INS-6
New Support for DNS/BIND as Primary Name Service	INS-6
NET\$STARTUP Procedure	

(continued on next page)

Documentation Errata and Addenda 4.2 Documentation Update for Version 8.2

Table 4–1 (Cont.) Documentation Update/Release Notes Cross-Reference

Release Note Topic	Refer to ...
NET\$STARTUP supports STARTUP_P2 SYSGEN procedure	MNG-8
NET\$LOGICALS File	
New logicals	MNG-6
NET\$SHUTDOWN	
New NET\$APPLICATION_SHUTDOWN logical	MNG-6, MNG-8
NCL	
Explanation of IP address backtranslations	NCL-1
New management topics in NCL Help	NCL-1
Improved explanation of access control	NCL-1
Improved explanation of wildcard proxies	NCL-1
New SHUT ALIAS PORT command	NCL-3
Host Based Routing	
Routing Originating Queue Limit characteristic modified	NCL-21
Session Control/CDI Cache	
CDI cache improvement	MNG-5
CDI\$SYSTEM_TABLE logical name table	MNG-5
Resolution of IP fully-qualified names	MNG-5
DECdns	
Establishing clerk connections to WAN TCP/IP Servers	INS-7
New logicals for the DNS\$ADVER process	DNS-12
Controlling the LAN devices used by DECdts	DNS-6
DECdts	
New procedures for Changing Time	INS-7 ¹
New DECdts NTP provider sample program	DTS-C
Designating Global Servers Using the BIND Database	DTS-2
Option to Disable DECdts at System Boot	DTS-2
Controlling the LAN devices used by DECdts	DTS-2
OSAK	
<i>OSAK new features moved to Documentation Errata chapter.</i>	
Documentation Corrections	
NCL	
Syntax corrections	NCL-(throughout command descriptions)

¹See also the "Setting System Time" chapter in the *HP OpenVMS System Manager's Utility Guide*, the "Creating Time Zone Conversion Information" in the "Overview" chapter of the *Compaq C Run-Time Library Utilities Reference Manual*, and the "Date/Time Functions" chapter and tzset function description in the *HP C Run-Time Library Reference Manual for OpenVMS Systems manual*.

(continued on next page)

Documentation Errata and Addenda

4.2 Documentation Update for Version 8.2

Table 4–1 (Cont.) Documentation Update/Release Notes Cross-Reference

Release Note Topic	Refer to ...
Alias Port Outgoing Default characteristic added	NCL-3
Routing Entity Maximum Buffers characteristic modified	NCL-21
Removed restriction on modifying OSI Transport Cons Filters characteristic	NCL-20
Added material about using SET NCL DEFAULT commands	NCL-1
Added material about using NCL SNAPSHOT command	NCL-1
Network Management	
New DECNET_VERSION values	MNG-3
Added material about using SNAPSHOT command	MNG-11, NCL-1
Operational Notes	
Network Management	
Startup procedure changed: NCL information minimized	MNG-6, MNG-8
decnet_register option 3 - updating registered node towers	MNG-5, MNG-E
DECdns	
Modifying a Clerk to use the Cluster Alias	DNS-6
Providing adequate resources for DNS clerk	DNS-12
Providing adequate resources for DNS server	DNS-12
Using the DNS.CONF file	DNS-12
Adjusting null_port_timeout in DNS.CONF file	DNS-12
Note on using the NEW EPOCH command	DNS-11
Default parameters for process limits on DNS servers	DNS-12
Removing Obsolete DNS\$CACHE files	DNS-12

4.3 Network Management and NCL

The following sections describe documentation errata and addenda for the *HP DECnet-Plus for OpenVMS Network Management*, *HP DECnet-Plus for OpenVMS Network Control Language Reference*, and *DECnet-Plus Problem Solving* manuals.

4.3.1 Devices and OpenVMS Device Names

Version 8.2 The sections that follow list the devices that can be configured using DECnet-Plus software. The last column lists the device name to use for the communication port characteristic when creating entities that reference these devices.

Note

DECnet-Plus supports a subset of the communications devices supported by OpenVMS. Some of the devices shown in these tables may be configurable (either by NET\$CONFIGURE or directly by NCL) but are no longer supported by DECnet-Plus. See the *Software Product*

Descriptions (SPDs) for HP DECnet-Plus for OpenVMS and HP OpenVMS for a list of devices supported by the current release.

4.3.1.1 CSMA-CD Devices

The circuit devices listed in Table 4–2 can be configured using DECnet-Plus software. They provide multi-access connections among many nodes on the same CSMA-CD circuit.

Table 4–2 CSMA-CD Devices

Device Name	Type	OS	I/O Bus	OpenVMS Name	DECnet Comm Port
82559 (Intel)	100BaseTX	I64	PCI(embedded)	EI	EIc
A5230A	100BaseTX	I64	PCI	EW	EWc
A5506B	100BaseTX (4-port)	I64	PCI	EW	EWc
A6825A	1000BaseTX	I64	PCI	EW	EWc
A6847A	1000BaseSX	I64	PCI	EW	EWc
A7011A	1000BaseSX (2-port)	I64	PCI-X	EI	EIc
A7012A	1000BaseTX (2-port)	I64	PCI-X	EI	EIc
A9782A	1000BaseSX	I64	PCI-X	EW	EWc
A9784A	1000BaseTX	I64	PCI-X	EW	EWc
BCM5701	1000BaseTX	I64,Alpha	PCI(embedded)	EW	EWc
BCM5703	1000BaseTX	I64,Alpha	PCI(embedded)	EW	EWc
COREIO	10Base5	Alpha	TURBOchannel	ES	ESc
DEBNA	10Base5	VAX	BI-bus	ET	ETc
DEBNI	10Base5	VAX	BI-bus	ET	ETc
DEBNK	10Base5	VAX	BI-bus	ET	ETc
DEBNT	10Base5	VAX	BI-bus	ET	ETc
DEGPA-xx	1000BaseTX,SX	Alpha	PCI	EW	EWc
DEGXA-xx	1000BaseTX,SX	Alpha	PCI,PCI-X	EW	EWc
DELQA	10Base5	VAX	Q-bus	XQ	XQc
DELTA	10Base5 (2-port)	VAX	TURBOchannel	EC	ECc
DELUA	10Base5	VAX	UNIBUS	XE	XEc
DEMNA	10Base5	VAX,Alpha	XMI	EX	EXc
DEQTA	10Base5	VAX	Q-bus	XQ	XQc
DESQA	10Base5	VAX	Q-bus	ES	ESc
DEVA	10Base5	VAX	embedded	ES	ESc
DEUNA	10Base5	VAX	UNIBUS	XE	XEc
DE20c	10Base2,5,T	Alpha	ISA	ER	ERc
DE422	10Base2/T	Alpha	EISA	ER	ERc
DE425	10Base2/5/T	Alpha	EISA	ER	ERc
DE434	10BaseT	Alpha	PCI	EW	EWc

(continued on next page)

Documentation Errata and Addenda

4.3 Network Management and NCL

Table 4–2 (Cont.) CSMA-CD Devices

Device Name	Type	OS	I/O Bus	OpenVMS Name	DECnet Comm Port
DE435	10Base2/5/T	Alpha	PCI	EW	EWc
DE436	10BaseT (4-port)	Alpha	PCI	EW	EWc
DE450	10Base2/5/T	Alpha	PCI	EW	EWc
DE500-xx	100BaseTX,FX	Alpha	PCI	EW	EWc
DE504-BA	100BaseTX (4-port)	Alpha	PCI	EW	EWc
DE600-xx	100BaseTX,FX	Alpha	PCI	EI	EIc
DE602-xx	100BaseTX,FX (2-port)	Alpha	PCI	EI	EIc
KFE52	10Base5	VAX	Q-bus	EF	EFc
P2SE	100BaseTX	Alpha	PCI	EW	EWc
P2SE+	100BaseTX	Alpha	PCI	EW	EWc
PMAD	10Base5	VAX,Alpha	TURBOchannel	EC	ECc
SGEC	10Base2	VAX	Local	EZ	EZc
TGEC	10Base2	Alpha	Local	EZ	EZc

Note

ATM devices are supported in emulated LAN (ELAN) mode by the OpenVMS EL device. For information about configuring an ELAN (and creating the EL device), see the "Configuring an Emulated LAN (ELAN)" section in the "Local Area Network (LAN) Device Drivers" chapter in the *HP OpenVMS I/O Users Reference Manual*.

LAN failover sets are supported by the OpenVMS LL device. For a general discussion of LAN failover sets, see Chapter 10 in Volume 2 of the *HP OpenVMS System Manager's Manual*. For information about managing LAN failover sets, see the description of the LANCP utility's DEFINE DEVICE and SET DEVICE commands in the *HP OpenVMS System Management Utilities Reference Manual*.

4.3.1.2 FDDI Devices

The devices listed in Table 4–3 can be configured using DECnet-Plus software. These devices provide multi-access connections among many nodes on the same FDDI circuit.

Table 4–3 FDDI Devices

Device Name	Type	OS	I/O Bus	OpenVMS Name	DECnet Comm Port
DEFAA	100 mmf	Alpha	FutureBus+	FA	FAc
DEFEA-xx	100 mmf, UTP	Alpha	EISA	FR	FRc
DEFPA-xx	100 mmf, UTP	Alpha	PCI	FW	FWc
DEFPZ-xx	100 mmf, UTP	Alpha	PCI	FW	FWc

(continued on next page)

Table 4–3 (Cont.) FDDI Devices

Device Name	Type	OS	I/O Bus	OpenVMS Name	DECnet Comm Port
DEFQA	100 mmf	VAX	Q-bus	FQ	FQc
DEFTA-xx	100 mmf, UTP	VAX,Alpha	TURBOchannel	FC	FCc
DEFZA	100 mmf	VAX,Alpha	TURBOchannel	FC	FCc
DEMFA	100 mmf	VAX,Alpha	XMI	FX	FXc

4.3.1.3 Synchronous Devices

The synchronous devices listed in Table 4–4 can be configured by DECnet-Plus software. All of the synchronous line devices are either point-to-point or multipoint tributary circuit devices.

Table 4–4 Synchronous Devices

Device Name	Type	OS	I/O Bus	OpenVMS Name	DECnet Comm Port
DIV32		VAX	Q-bus	SD	DIV-c-u
DMB32	1-port	VAX	VAXBI	SI	DMB-c-u
DMF32	1-port	VAX	UNIBUS	XG	DMF-c-u
DNSES	2-port	Alpha	EISA	ZE	ZE-c-u
DPV11	1-port	VAX	Q-bus	SE	DPV-c-u
DSB32	2-port	VAX	VAXBI	SL	DSB-c-u
DSF32	2-port	VAX	MI-bus	SF	DSF-c-u
DSH32	2-port	VAX	Busless	ZS	DSH-c-u
DST32	1-port	VAX	Busless	ZS	DST-c-u
DSV11	2-port	VAX	Q-bus	SJ	DSV-c-u
DSW21	1-port	VAX	Busless	ZT	DSW-c-u
DSW41	1-port	VAX	Busless	ZT	DSW-c-u
DSW42	2-port	VAX	Busless	ZT	DSW-c-u
DSYT1	2-port	Alpha	TURBOchannel	ZY	ZY-c-u
DUP11	1-port	VAX	UNIBUS	SE	DPV-c-u
PBXDD-Ax	2-, 4-port	Alpha	PCI	ZZ	ZZ-c-u
PBXDI-Ax	2-, 4-port	Alpha	ISA	ZS	ZS-c-u
PBXDP-Ax	2-, 4-, 8-port	Alpha	PCI	ZW	ZW-c-u
SCC	1-port	Alpha	None	ZR	ZR-c-u

Note

For DECnet-Plus for OpenVMS VAX systems, see the Software Product Description (SPD) for maximum supported line speeds for the various Data Link layer protocols. The SPD also contains processor loading information specific to X.25 communications.

Documentation Errata and Addenda

4.3 Network Management and NCL

4.3.1.4 Asynchronous Devices (OpenVMS VAX Only)

The asynchronous devices listed in Table 4–5 can be configured by DECnet-Plus software on OpenVMS VAX systems.

Table 4–5 Asynchronous Devices (OpenVMS VAX Only)

Device Name	Type†	I/O Bus	OpenVMS Name	DECnet Comm Port
DHQ11	8-line,MC	Q-bus	TX	TX-c-u
DHU11	16-line,MC	UNIBUS	TX	TX-c-u
DHV11	8-line,MC	Q-bus	TX	TX-c-u
DMB32	8-line,MC	VAXBI	TX	TX-c-u
DMF32	8-line,MC	UNIBUS	TX	TX-c-u
DMZ32	24-line,MC	UNIBUS	TX	TX-c-u
DZ11	8-line,PMC	UNIBUS	TT	TT-c-u
DZ32	8-line,PMC	UNIBUS	TT	TT-c-u
DZQ11	4-line,PMC	Q-bus	TT	TT-c-u
DZV11	4-line,PMC	Q-bus	TT	TT-c-u

†MC = full modem control, PMC = partial modem control

4.3.2 DTR/DTS

Version 7.3

In Section 3.10.2 of the *DECnet-Plus Problem Solving* manual, the documentation for the /nodename qualifier of the DTS test command contains a statement that seems to indicate that DTS/DTR might support Phase V node names. The statement should read as follows:

The *node-id* can be 0, a Phase IV address or node name, or, on a Tru64 UNIX system, a DECnet-Plus NSAP address or Phase V node name. To specify a Phase V node on OpenVMS systems, you must use a node synonym of 6 characters or less.

4.4 OSAK Programming

The following sections provide information that updates the current OSI Applications Kernel (OSAK) programming documentation.

4.4.1 called_aei Parameter on A-ASSOCIATE and S-CONNECT Indication Events

Version 7.2

For A-ASSOCIATE indication events in the OSAK API, and S-CONNECT indication events in the SPI, the *nsap* field of the *called_aei* parameter is no longer filled in by OSAK. Previously, the *called_aei.paddress.nsap* structure contained the values specified in the *local_aei.paddress.nsap* parameter on the *osak_open_responder* or *spi_open_responder* call. However, these values are not used by *osak_open_responder* or *spi_open_responder* and may not contain correct information about the received connection (for example, the transport type).

4.4.2 ROSE Minimum Workspace Increased

Version 7.2 The minimum required size for the parameter block workspace for ROSE has increased. If your ROSE application is using the old minimum workspace size, you will now get a bad parameter error. It is recommended that your application use the constant `ROSE_WS_SIZE` defined in `osak_api.h` to get the correct minimum workspace size.

Note

Depending on the size of your application's parameter block workspace, you may have to recompile your ROSE applications.

4.4.3 Using OSAK with OSI over TCP/IP

Version 7.2 On the initiator side, to make a connection using OSI over TCP/IP, specify as the NSAP the IP address as a 6-byte nibble packed hex value including the port number in the first two bytes. For example, the address of 16.36.112.142 on port 102 would be encoded as follows:

```
102  16.36.112.142
|    |    |    |
0066 10 24 70 8E
```

In this case, the NSAP would be 00661024708E (hex). Specify the NSAP type as `OSAK_C_RFC1006`. For the transport template, specify either the pseudo-template "1006" or no template at all. OSAK uses a default template of "1006" if the NSAP type is `OSAK_C_RFC1006`.

On the responder side, to listen for OSI over TCP/IP connections, specify as the transport template the pseudo-template "1006".

4.4.4 New Status Code for `osak_send_more` and `spl_send_more` Routines

Version 7.2 The routines `osak_send_more` and `spl_send_more` can now return a status code of `OSAK_S_NODATA` which indicates that there is no data remaining. In previous versions, the status code `OSAK_S_INVFUNC` may have been returned when this error occurred.

4.4.5 New OSAK API Routine: `osak_tc_no_sc()`

Version 7.2 OSAK now contains a new API routine which allows application programs to query OSAK to determine whether a transport connection has been established on a port, but a session connection has not.

Syntax:

```
status = osak_tc_no_sc(port, state)
```

C Binding:

```
osak_tc_no_sc(port, state)
osak_port port;
int *state;
```

Arguments:

port

The OSAK port that you wish to query.

state

Documentation Errata and Addenda

4.4 OSAK Programming

A pointer to an integer where OSAK should return the state information. If the port is connected at the transport layer, but not at the session layer, OSAK returns a non-zero value (true). If the port is connected at the session layer, OSAK returns a zero value (false).

Description:

Call this routine to determine whether the port is connected at the transport layer, but not at the session layer. This information can be useful when interoperating with some other OSI implementations that may incorrectly terminate the transport connection, or may establish a transport connection but never complete the session connection. Use of this routine can assist an application in determining if the OSAK port is hung. An equivalent routine also exists for the SPI interface: `spi_tc_no_sc()`.

Return Values:

OSAK_S_NORMAL The routine has finished without error.
OSAK_S_INVPORT The port identifier is invalid.

4.4.6 Handling of Default Transport Templates

Version 7.1 The following information clarifies statements made about template defaulting in the *DECnet-Plus OSAK Programming Reference* and *DECnet-Plus OSAK SPI Programming Reference* manuals.

If no transport template is specified in the `osak_associate_req` (API and ROSE) or `spi_connect_req` (SPI) call, OSAK takes the following actions:

- If the NSAP in the `osak_nsap` data structure is CONS or CLNS, OSAK uses the default transport template.
- If the NSAP in the `osak_nsap` data structure is RFC 1006, OSAK uses the `osit$rfc1006` transport template.

4.4.7 Omitted Parameter Block Parameters

Version 7.1 In the *DECnet-Plus OSAK SPI Programming Reference* manual, the following OSAK parameter block parameters were omitted from the SPI routine descriptions. Although the parameters are listed in Section 1.2 and described in Section 1.4.2, they should also appear in the Parameters Used table as follows:

Parameters Used	Data Type	Access
[user_context]	Address	read only
port_id	osak_port	write only

The preceding parameters should be listed for all SPI routines that have the OSAK parameter block argument, except for the following:

- `spi_async_close`
- `spi_close_port`
- `spi_collect_pb`
- `spi_expedited_req`

Operational Notes

This chapter contains operational notes and restrictions for HP DECnet-Plus for OpenVMS. This chapter also includes known problems with the current release.

5.1 Network Management

The following sections discuss several operational notes about network management.

5.1.1 CTF Utility Now Requires Process Privileges

Version 8.2

The CTF utility has always required that the invoking process have certain rights identifiers in order to use the START and STOP commands to control the collection of trace records. These rights identifiers are described in Chapter 3 of the *DECnet/OSI for VMS CTF Use* manual.

Now, in addition to one or more of the required rights identifiers, the invoking process must have the following process privileges before invoking CTF for any purpose:

ALTPRI, BYPASS, CMKRNL, NETMBX, PRMGBL, PRMMBX, PSWAPM, SYSGBL, SYSLCK, SYSNAM, TMPMBX, and WORLD

If you invoke CTF from a process without the required process privileges, the utility exits with an error message indicating the first of the required privileges not granted.

5.1.2 Tuning to Support the Larger CML BYTLM Value

*Version 8.2
and 7.3-1*

The default BYTLM value for CML has been increased from 10,000 to 64,000. This correction was done to prevent NCL from returning the NCL-E-OVERFLOWBYTLM error message when large amounts of data were needed from CML. If you have ever experienced this problem, or if your MAXBUF value is already set to 64000 and you have received NCL-E-NORESOURCEAVAIL error messages, then you should make the changes described in this section.

To fully take advantage of the increased BYTLM value, do the following:

1. Increase the value of MAXBUF to 64,000.
2. Raise the BYTLM quota on any account that might issue NCL commands requiring the increased resources to at least 65,000. Be sure to do this to at least the SYSTEM, CML\$SERVER, and MIRRO\$SERVER accounts. For example, to modify the CML\$SERVER account, use the following command:

```
$ run sys$system:authorize
UAF> modify cml$server/bytlim=65000
```

Operational Notes

5.1 Network Management

3. Check for the existence of the logicals NET\$ACP_BYTLM, NET\$EVD_BYTLIM, and NET\$MOP_BYTLIM. If any of these logicals are defined and they are set to a value lower than 65,000, modify their entries in SYS\$MANAGER:NET\$LOGICALS.COM to raise their value to at least 65,000.
4. Check for the existence of the NET\$LESACP_BUFFER_LIMIT logical. If this logical is defined and it is set to a value lower than 65,000, modify its entry in SYS\$STARTUP:NET\$LES_STARTUP.COM to raise the value to at least 65,000.
5. Check for the existence of the NET\$ADVER_BUFFER_LIMIT logical. If this logical is defined and it is set to a value lower than 300,000, modify its entry in SYS\$STARTUP:SYLOGICALS.COM to raise the value to at least 300,000.
6. Reboot the system to make all the preceding changes effective.

If, after applying the measures described in the previous list, you now receive NCL-E-NORESOURCEAVAIL errors instead of NCL-E-OVERFLOWBYTLIM errors, this generally indicates that your process BYTLM quota is not set high enough. Set the quota to a value slightly higher than the MAXBUF value. Setting process BYTLM quotas to 65,000 (as specified in the previous list) should be sufficient to avoid this problem.

5.1.3 Upgrading to DECnet-Plus with Systems Running DEC X.25 Client Software (Alpha Only)

Version 7.1 If you are upgrading to DECnet-Plus from DECnet Phase IV and have the DEC X.25 Client for OpenVMS Alpha Systems product installed, be aware that there are two X.25 products for OpenVMS Alpha systems:

- X.25 client software, for DECnet Phase IV systems only
- X.25 software, for DECnet-Plus systems only

If you upgrade to DECnet-Plus and want to retain access to X.25, remove the DEC X.25 Client product and install the X.25 product, which provides both client and native X.25 functionality.

For details about installing and configuring the X.25 product, refer to the *HP DECnet-Plus for OpenVMS Installation and Configuration* and the *HP X.25 for OpenVMS Configuration* manuals.

5.1.4 COLLECT Command and Addressing in decnet_migrate Commands

Version 7.1 Some connection errors might be reported for nodes with DECnet-Plus addresses that are not Phase-IV compatible as well. These errors occur because the network management interface on the OpenVMS operating system does not recognize connections to nodes by the explicit use of a DECnet-Plus address that is not Phase-IV compatible.

If a DECnet-Plus node also has a Phase-IV compatible address, information is collected for it using that address.

5.1.5 Decimal Syntax DSPs Not Supported

Version 7.1 DECnet-Plus supports only those IDPs (initial domain parts) that specify the use of binary syntax DSPs (domain-specific parts). IDPs that specify decimal syntax DSPs are not supported.

5.2 Devices and Data Links

The following sections discuss operational notes about device support and data links.

5.2.1 Receive Buffers Attribute on HDLC LINK Entity Not Implemented (I64 and Alpha Only)

Version 7.1 The receive buffers attribute of the `hdlc link` entity is not implemented in DECnet-Plus. If you attempt to set the attribute according to the following example, NCL displays the subsequent informational messages:

```
NCL>set node 0 hdlc link hdlc-0 receive buffers 8
```

```
Node 0 HDLC Link HDLC-0  
at 2004-04-28-09:51:10.780-04:00I0.113
```

```
Command failed due to:  
set list error
```

```
Characteristics  
no such attribute ID: Receive Buffers = 8
```

HDLC (High-level Data Link Control) allocates the number of receive buffers that it requires; you cannot specify this setting.

5.2.2 FDDI Large Packet Support

Version 7.1 Phase IV nodes do not support FDDI large packets. If your node has an FDDI controller that uses large packets, routing selects a smaller segment size for Phase IV adjacencies. To communicate with Phase IV nodes, you must enable Phase IV addressing on the routing circuit. Verify that the `enable phaseiv address` attribute is set to true on one FDDI circuit in the `SY$MANAGER:NET$ROUTING_STARTUP.NCL` file.

5.3 Transport Layer

The following sections discuss operational notes for the transports in DECnet-Plus.

5.3.1 NSP Characteristics for X.25 PVCs

Version 7.3 If a node has routing circuits of type `x25 permanent` you should set the NSP flow control policy characteristic to no flow control. In addition, if the call packet window size is 8 or less, you should set the NSP maximum window characteristic to 8. This lowers the number of NSP unacknowledged transmits and thus prevents the NSP out-of-order cache from filling.

5.3.2 NSP Flow Control Policy and DECnis and x25Gateway Routers

Version 8.2 In DECnet-Plus Version 6.3, several corrections were made to NSP flow control algorithms. Previous to these changes, if segment flow control was used flow control credits were sometimes improperly handled and the connections would appear to hang. The NSP software in DECnis routers and the x25Gateway product is based on the pre-DECnet-Plus Version 6.3 software. For this reason, whenever a connection is set up to one of these hardware products the NSP flow control should always be set to no flow control. Note that segment flow control is the default for NSP's flow control policy characteristic. Therefore, you must explicitly select NONE when running `NET$CONFIGURE`, use NCL to modify the running system, or edit the NCL startup script created by `NET$CONFIGURE`.

Operational Notes

5.3 Transport Layer

5.3.3 OSI Transport

Version 7.2 The following notes apply to OSI transport:

- If OSI transport times out during data transfer mode because the remote end is unreachable, the error `SS$_CONNECFAIL` may also be returned in the IOSB for the `$QIO(IO$_READVBLK)` or `$QIO(IO$_WRITEVBLK)` call. User code should handle `SS$_TIMEOUT` as well as `SS$_CONNECFAIL`.
- When defining an RFC 1006 address using the `OSIT$NAMES` logical name table, you must put a quotation mark around the address portion.
- The requirement against using 0 (zero) and 1 as valid NSELs (network selectors) has been removed, but the following restrictions apply:
 - The selector fields for OSI transport must be the same in both transport partners if the packet is traversing a backbone containing Phase IV routers (as opposed to routing vector domains).
 - Although the cluster alias uses two ranges of network transport selector fields to reduce the risk of a collision, HP recommends that you do not choose values from those ranges, or at least that you do not choose the first few values in each range.
- The `osi transport` entity does not support a value of 0 for the maximum remote NSAPs (network service access points) attribute.
- The `osi transport template` entity does not support a value of any for the network service attribute. If this attribute is set to any, it is treated by OSI transport as `clns`.

5.4 Session Control

The following sections discuss operational notes about Session Control.

5.4.1 Transport Precedence in Predominantly Phase IV Networks

Version 7.3 When using a Phase V node in a predominantly Phase IV network, the default for session control's transport precedence characteristic is `{OSI,NSP}`. This causes session control to attempt an OSI connection first, followed by an NSP connection. If you have a predominantly Phase IV network, this causes an unnecessary delay when connecting to Phase IV nodes. If the precedence is causing the delay, the delay is not seen when connecting in the reverse. To change the precedence, insert the following line at the end of the `NET$SESSION_STARTUP.NCL` file:

```
set node 0 session control transport precedence = {NSP,OSI}
```

Remember to reinsert this line if you run `NET$CONFIGURE`.

5.4.2 Session Control Back Translation Soft Links

Version 7.3–2 Although the default for Session Control's maintain backward soft links characteristic is correctly documented as `false`, on OpenVMS systems the Session Control startup script always sets this characteristic to `true`.

5.4.3 Restriction When Using Applications Supported Using Cluster Aliases

Version 8.2 Due to a design restriction, incoming ALIAS connections are forwarded to cluster members somewhat differently depending upon whether the SESSION CONTROL APPLICATION entity's ADDRESS set contains a NUMBER parameter (equivalent to Phase IV's connect by object number) or a TASK parameter (equivalent to Phase IV's connect by object name). For an incoming connect by object number, each ALIAS entity has access to a cluster-wide object number table that specifies which object numbers have been defined for each cluster member and the status of their INCOMING ALIAS characteristic. The result is that the connection forwarding mechanism works as expected, that is, connections are forwarded or not forwarded to cluster members correctly depending upon whether the application exists on the member and depending on the setting of the INCOMING ALIAS parameter.

However, the mechanism for mapping incoming connections by object name using this cluster-wide object number table does not exist and the result is that connections are forwarded to a given cluster member regardless of whether the application exists on a given member and regardless of the setting of the INCOMING ALIAS parameter on a given member. If you wish to restrict access to a network application supported using a cluster alias (either by not having the application present on a member or by setting the INCOMING ALIAS characteristic to FALSE on a member), you must create the SESSION CONTROL APPLICATION by supplying the NUMBER keyword for the ADDRESS characteristic (that is, supplying the object number), so that connection forwarding works as expected.

5.5 Distributed Naming Service (DECdns)

The following sections discuss operational notes for DECdns.

5.5.1 Creating Clearinghouses When IP Address is Zero Causes DNS Server Hang

Version 7.3-2 Do not create a clearinghouse when the node's IP address is all zeros. You can determine if the address is zero by using the NCL command SHOW ADDRESS.

After each system boot, the local node's IP address is zero until the first incoming or outgoing OSI over TCP/IP or DECnet over TCP/IP connection occurs on the local node. At that time, the PATHWORKS Internet Protocol (PWIP) driver included with HP TCP/IP Services for OpenVMS provides DECnet-Plus with its IP address. Consequently, at least one OSI over TCP/IP or DECnet over TCP/IP connection must occur before attempting to create a clearinghouse.

An easy method to ensure that at least one OSI over TCP/IP or DECnet over TCP/IP connection has been made is to use the SET HOST command with the IP address of the local node. (If you don't know the local node's IP address, use the DCL SHOW LOGICAL *HOSTADDR command to display it.) The SET HOST command causes HP TCP/IP Services for OpenVMS to provide the local IP address to DECnet-Plus. Remember to use the NCL SHOW ADDRESS command to verify that the local node's IP address has been defined.

If fail to do this, you may have to delete the clearinghouse and start over after you have obtained a non-zero IP address.

If you try to use a clearinghouse with zero IP addresses in the towers it is likely that the DECdns server will hang.

Operational Notes

5.5 Distributed Naming Service (DECdns)

5.5.2 Server Startup Delay

Version 7.3-2 There can be a delay in the startup of the DECdns server when the server is using DECnet over TCP/IP connections. This delay is due to the time required for the server to obtain a non-zero IP address from DECnet and the PATHWORKS Internet Protocol (PWIP) software. The IP address is initially zero until it is updated by DECnet and PWIP with the actual IP address for the node.

5.5.3 Delay in Node Name-to-Address Translations When DECdns Is Using DECnet over TCP/IP Connections

Version 7.3-2 When DECdns is using DECnet over TCP/IP connections, you may experience delays whenever DECnet has to translate a node name to a node address using the node's IP tower. These delays can occur when you are using NCL, when you invoke a DECnet application, or when you perform any other action that requires translating the node's name to its IP address.

The DECdns clerk, like any other DECnet application, uses DECnet to access the DECdns server. Currently, Session Control's transport precedence characteristic does not support a value for the IP tower. Therefore, Session Control always attempts to connect to the DECdns server using the NSP and OSI towers before trying the IP tower. After DECnet notifies DECdns that it cannot make a connection on a OSI or NSP tower, the DECdns clerk tries the IP tower.

Note

This problem is especially noticeable in environments where a process spawns multiple subprocesses, each of which performs a network command. This is common in some batch environments.

5.5.4 Problem Creating Clearinghouse Using DECnet over TCP/IP Connections

Version 7.3-2 When creating a new clearinghouse that joins an existing namespace in an environment that uses DECnet over TCP/IP connections, it is possible to have a clearinghouse created without the root directory. This is usually due to an intermittent communications error during clearinghouse creation. This problem occurs infrequently.

The current workaround is to force the creation of the root directory by issuing a manual skulk command as shown below:

```
$ mcr dns$control
DNS>set dir . to skulk
```

5.5.5 DNS\$CONFIGURE Option 4 Does Not Display IP Address

Version 7.3-2 Option 4 of the DECdns configuration program (DNS\$CONFIGURE.COM) is used to display a server node's address (as described in Section 10.5 of the *HP DECnet-Plus for OpenVMS DECdns Management* manual). This option has not been enhanced to display the node's IP address. As a workaround, use the NCL show node [*node-id*] address command.

5.5.6 Limitation on Size of Checkpoint File Records

Version 7.3 The checkpoint file has a limitation of 500 blocks maximum on the size of individual records. To detect any records over this limit or approaching this limit, run the SURGEON "-scanrx" option. You will need to take corrective action if you see a line indicating greater than 500 blocks and the number of records is one (RECS: 1) because the server may bugcheck with a 902 error in the DB_COMPRESSION.C module.

Records that exceed the 500 block maximum can be caused by an application that is writing records larger than this limit or the garbage collector mechanism (removes deleted records) is not running.

To correct the problem, apply one of the following actions:

- Use the SURGEON "-exciser" option to remove the offending records
- Use the recreate object command (if the garbage collector is not running).

If you have any questions about doing this, consult your HP representative.

Avoid the situation with the garbage collector not running by making sure all your directories skulk to a success status and that the DNS\$ALLUPTO date is recent.

Use the following command in DNS\$CONTROL to check the skulk status of all directories in the namespace:

```
SHOW DIR ... DNS$SKULKSTATUS
```

Use the following command in DNS\$CONTROL to check the DNS\$ALLUPTO date on all directories in the namespace:

```
SHOW DIR ... DNS$ALLUPTO
```

5.5.7 PAGEDYN Resources: Clerks and Servers

Version 7.1 DECdns clerk and server systems can require additional paged dynamic memory resources (PAGEDYN).

For DECdns clerk systems, consider increasing the PAGEDYN resource on the node if you see either the RESOURCEERROR or NONSRESOURCES error during skulk operations.

Before you configure a system as a DECdns server, check that the system has at least 50,000 free bytes of paged dynamic memory. Insufficient paged dynamic memory on servers can cause configuration errors, skulk failures, and in some cases normal clerk operations can fail, returning either the RESOURCEERROR or NONSRESOURCES error.

The DNS\$SERVER.LOG file can also contain messages regarding failures caused by this resource.

Servers holding master replicas of directories that have many read-only replicas can also require additional paged dynamic memory.

Operational Notes

5.5 Distributed Naming Service (DECdns)

5.5.8 Error Message Generated by DECdns Clerk During Startup

Version 7.1 The DECdns clerk startup might output the following error message:

```
create node 0 dns clerk known namespace CZ
command failed due to:
process failure

A Known Namespace with this name or NSCTS already exists
```

You can disregard this message. The DECdns clerk configuration procedure puts a `create dns clerk known namespace` command for the default namespace into the DECdns clerk NCL startup file. The command remains there in case the DECdns cache file has been deleted or corrupted.

5.5.9 Miscellaneous DECdns Operational Notes

Version 7.1 You might encounter the following DECdns problems:

- DNS\$CONTROL might return Syntax Error if the clerk is disabled.
- You should not use DNS\$CONTROL to modify the `DNA_NodeSynonym` attribute on DECnet node objects. DNS\$CONTROL displays the attribute properly but does not modify it properly. Use the `decnet_register` node registration tool to modify DECnet data stored in DECdns.
- Re-created entries retain the case of the original entry name.
- If you have trouble creating a clearinghouse, make sure that both the `dns$server` account (user authorization file entry) and its default directory exist.

5.6 Distributed Time Service (DECdts)

The following sections discuss operational notes for DECdts.

5.6.1 Automatic Time Zone Changes on Rebooting Clusters

Version 7.1 If all members of a cluster are down when daylight saving time (DST) takes effect and automatic time zone changes are enabled, then members may reboot with the incorrect local time. If you anticipate that the cluster will be down during the change to or from daylight saving time, you should disable automatic time zone changes and make the changes manually.

5.7 OSAK Programming

The following sections discuss operational notes for the OSI Applications Kernel (OSAK) interface.

5.7.1 User Data Size Restriction for Tracing

Version 7.2 OSAK does not support tracing with user data buffers larger than 32,765 bytes.

5.7.2 OSAK Version 1.1 No Longer Supplied

Version 7.1 HP no longer provides the OSAK Version 1.1 interface.

If you use the OSAK Version 1.1 interface, you must migrate to the OSAK Version 3.0 SPI (session programming interface). The OSAK Version 3.0 SPI provides the same capabilities as the OSAK Version 1.1 interface.

For detailed information about the OSAK Version 3.0 SPI, see the *DECnet-Plus OSAK SPI Programming Reference* manual.

5.7.3 Migrating Applications from Version 2.0 to Version 3.0 OSAK API Interface

Version 7.1 Version 3.0 of OSAK includes a programming interface that is backwards compatible with the Version 2.0 programming interface. As a result of more rigorous conformance testing, some additional problems have been uncovered in the Version 2.0 interface.

As a result, applications built against OSAK Version 2.0 might encounter problems when using the Version 2.0 backwards-compatible interface in OSAK Version 3.0 if the applications contain previously untrapped application problems.

5.7.4 Misleading OSAK Error Message: SPM Congestion at Connect Time

Version 7.1 When OSAKserver receives an association request that specifies a nonexistent user identity, it generates the error message "SPM Congestion at Connect Time."

If you see this error, check that the user name and password specified in the request are correct. If the user name is correct but the password is incorrect, the OPCOM message %LOGIN-F-INVPWD is logged. If the user name is incorrect, no OPCOM message is logged.

5.8 DECnet Programming

The following sections discuss programming issues that apply to the DECnet programming interface.

5.8.1 64-Bit Virtual Address Restrictions

Version 7.1 You can use 64-bit virtual addresses only with \$QIO operations on the NET: devices; you cannot use 64-bit virtual addresses with \$QIO operations on the OS: devices.

For transparent task-to-task communication, the \$ASSIGN system service accepts only 32-bit addresses and 32-bit descriptors for the DEVNAM parameter when it describes a network connect block (NCB). You can use non-transparent task-to-task operations when you want 64-bit descriptor addresses and 64-bit descriptors to specify an NCB.

5.8.2 Backtranslation of IP Addresses

Version 7.1 For IP addresses, session control returns the full name if the synonym is not available regardless of the Session Control Application Node Synonym setting. If neither the full name nor synonym is available, the IP\$aa.bb.cc.dd format is returned.

If the system is configured to use IP addresses, you should prepare the application to receive a node name or address string longer than six characters in length even if Session Control Application Node Synonym = TRUE is specified for the application.

5.8.3 Applications Connection Requirement

Version 7.1 With DECnet-Plus, your application must wait for the completion of the \$QIO system service call used to accept the connection before attempting to use the connection.

