

OpenVMS Networking Updates



Prashanth K E
OpenVMS Engineering

Agenda

- TCP/IP Services V5.7 for OpenVMS update
- NFS enhancements
- CIFS enhancements
- LAN update
- DECnet Plus update
- Q & A

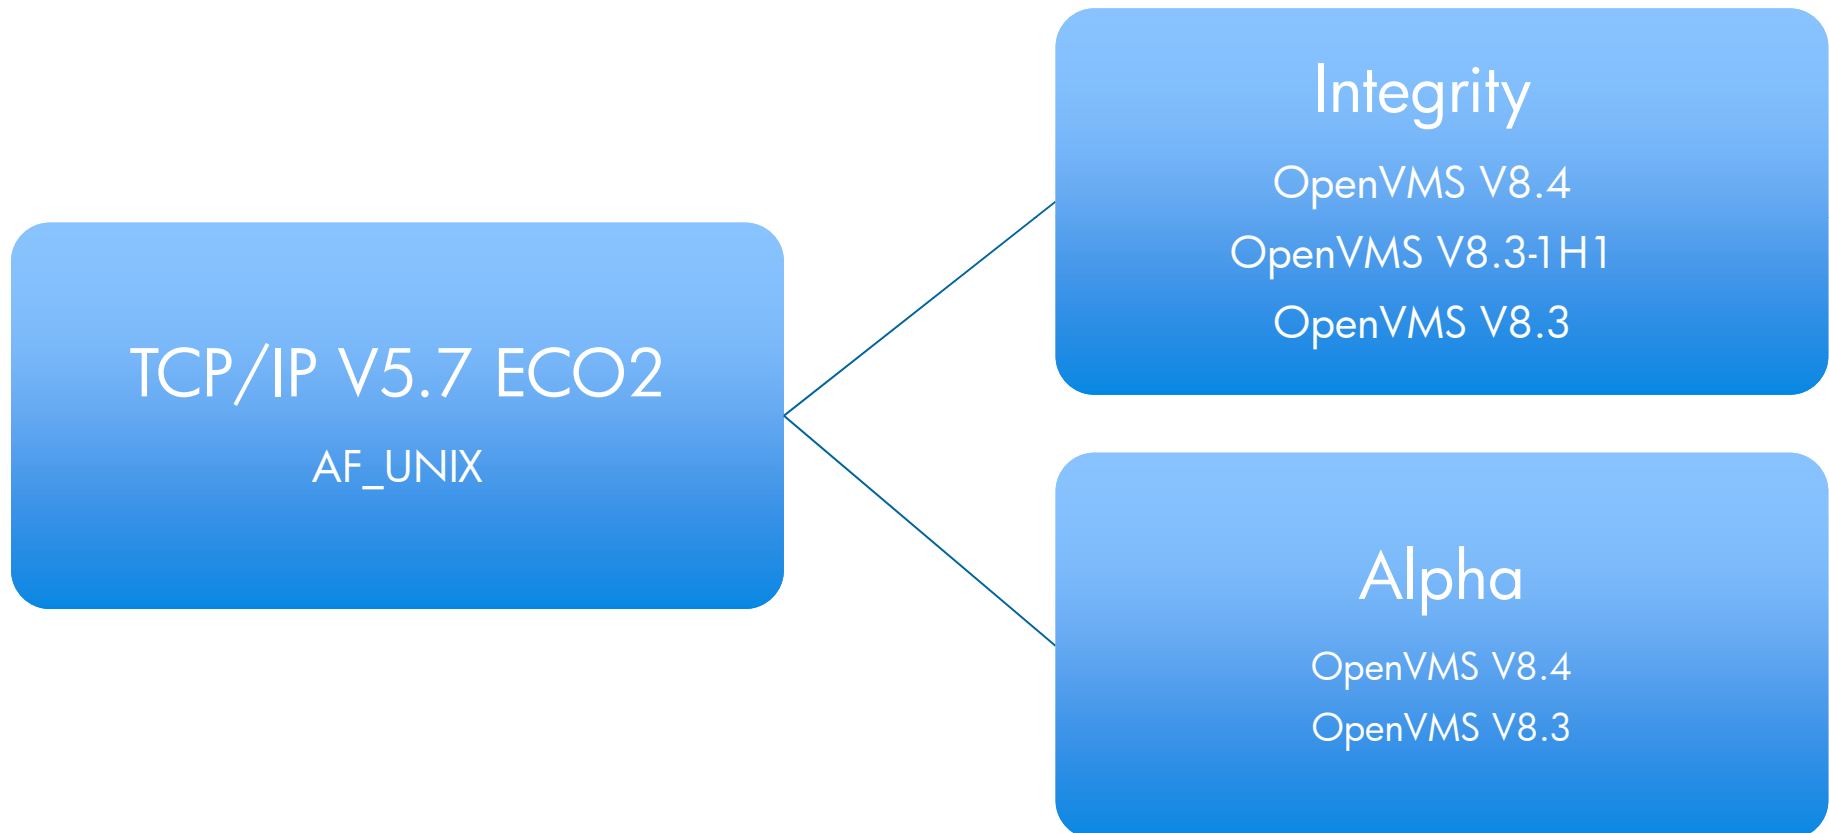


TCP/IP Services V5.7 for OpenVMS update



HP TCP/IP Services V5.7

Support Matrix



HP TCP/IP Services V5.7

Enhancements

•Cluster Over IP Enablement



•Packet Processing Engine



•Stream Control Transmission Protocol



•Enhanced FTP



•LPD Enhancements



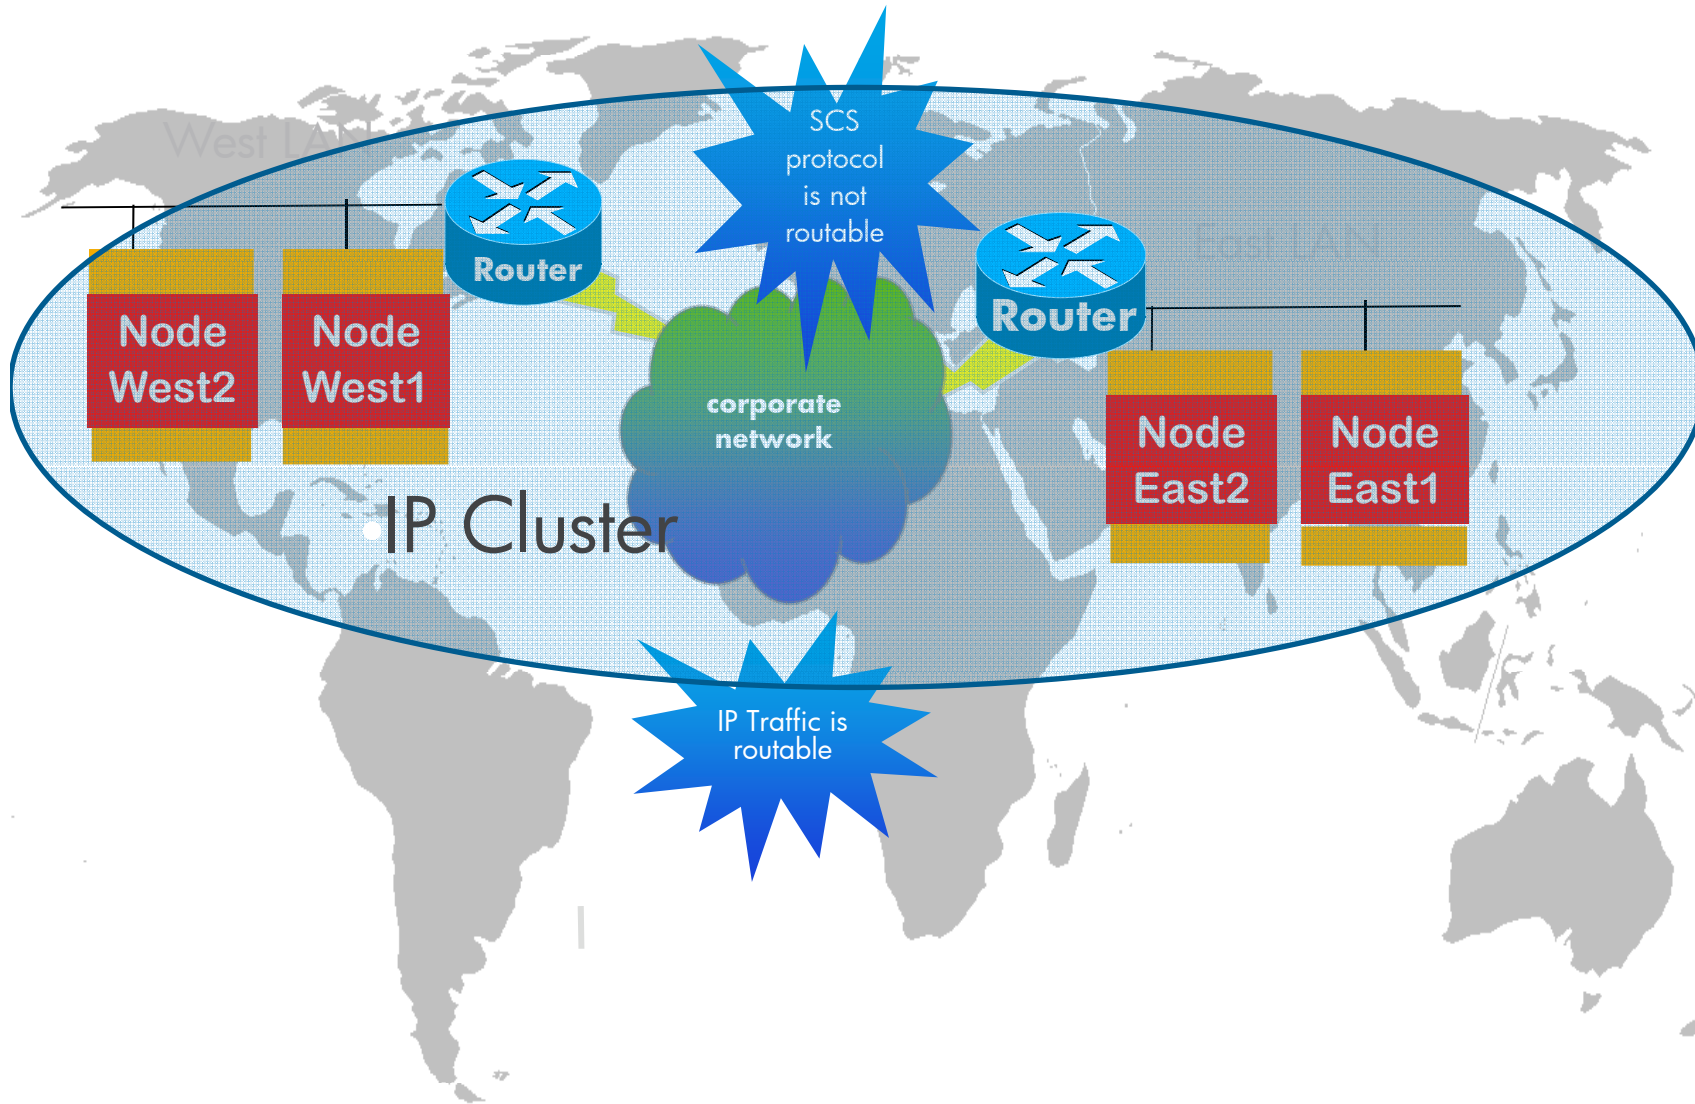
•Mail Enhancements



CLUSTER OVER IP ENABLEMENT



Why Cluster Over IP



IPCI Benefits

LAN Bridging

- No dependency on Vendors for LAN bridging
- No extra license/cost for LAN bridging (Layer 2 service)

IP Benefits

- Lower infrastructural and operational costs
- Leverage the benefits from the improvements in IP and LAN interconnect technology

Ability to co-exist in a modern datacenter without special setting



Cluster Over IP

TCP/IP Support

- OpenVMS V8.4 and TCP/IP V5.7
 - IPv4 only
 - Requires static IP addresses and IP Unicast
 - Coexists with LAN interconnect for Cluster communication

TCP/IP changes

- Boot time loading and Initialization
 - Existing boot sequence; LAN, PE driver, TCP/IP
 - Boot Sequence with IPCI ; LAN, TCP/IP, PE driver
 - Ability to make use of boot time configuration information.

IPCI configuration

- Ability to modify permanent database of other cluster Members
- Support for IPCI addresses



IPCI Configuration using TCPIP

Support for IPCI addresses

– Creates `sys$system:TCPIP$cluster.dat`

– Simple text file: do not edit manually

– Same as any address – but some special care

- IPCI addresses are accessed early during boot

- Hence the need for a simple text file

- Deleting an IPCI address may cause cluster to lose

quorum

- Shutting down TCP/IP with IPCI may cause loss of quorum

http://h71000.www7.hp.com/doc/84final/4477/4477pro_018.html



PACKET PROCESSING ENGINE -- PPE



Packet Processing Engine

Dedicates a CPU for TCP/IP processing

Without PPE

- TCP/IP runs on a standard CPU
- CPU reaches saturation, TCP/IP may become bottleneck

With PPE

- No more sharing
- PPE maybe enabled/disabled dynamically

Modeled on the OpenVMS Dedicated Lock Manager



Packet Processing Engine

Configuration Requirements

Hardware

- Requires more than one active CPU
- Will become dormant if only one CPU is left active
- Wake-up when more than one CPU becomes available
- Better suited to systems with many CPUs

Software

- Configure the BGO: as the only driver on the nominated CPU
- i.e. all other fast-path drivers must be moved to other CPUs



Example

```
$ show fastpath
```

```
Fast Path preferred CPUs on TEST 1-OCT-2010 10:25:10.80
```

```
HP BL860c (1.59GHz/9.0MB) with 4 active CPUs
```

```
Device:          Fastpath CPU:
```

```
EWA0             0
```

```
EWC0             2
```

```
FGA0             0
```

```
FGB0             3
```

```
PKA0             1
```

```
OpenVMS TCP/IP is currently running on CPU 1
```

```
$set dev FGB0/pref=1
```

```
$set dev bg0/pref=3
```



Example

```
$ show fastpath      ! ensure no other drivers on CPU3
```

```
Fast Path preferred CPUs on TEST 1-OCT-2010
```

```
10:27:05.64
```

```
HP BL860c (1.59GHz/9.0MB) with 4 active CPUs
```

```
Device:           Fastpath CPU:
```

```
EWA0              0
```

```
EWC0              2
```

```
FGA0              0
```

```
FGB0              1
```

```
PKA0              1
```

```
OpenVMS TCP/IP is currently running on CPU 3
```



Packet Processing Engine

Dynamically

- `sysconfig -r inet ppe_enable=1 ! 0` to disable
- `sysconfig -q inet ppe_enable`

Permanently

- `TCPIP$ETC:SYSCONFIGTAB.DAT`

TCPIP\$INETPPE process

- Started automatically when PPE is enabled
- Runs at Priority 63
- Hibernates when PPE is disabled



STREAM CONTROL TRANSMISSION PROTOCOL (SCTP)



SCTP Concepts

- Features of both TCP and UDP
- Association and streams
 - Address Head-Of-Line blocking
- Multihoming
 - Network resilience
 - Dynamic reconfiguration of IP addresses
 - Mix of v4 and v6 addresses
- Partial Reliability (How much reliability)
 - When should I retry transmission and when do I give up ?
- Security
 - 32 bit verification tag (vtag) exchanged during association setup
 - Needed to send / receive packets



Comparison – SCTP, TCP and UDP

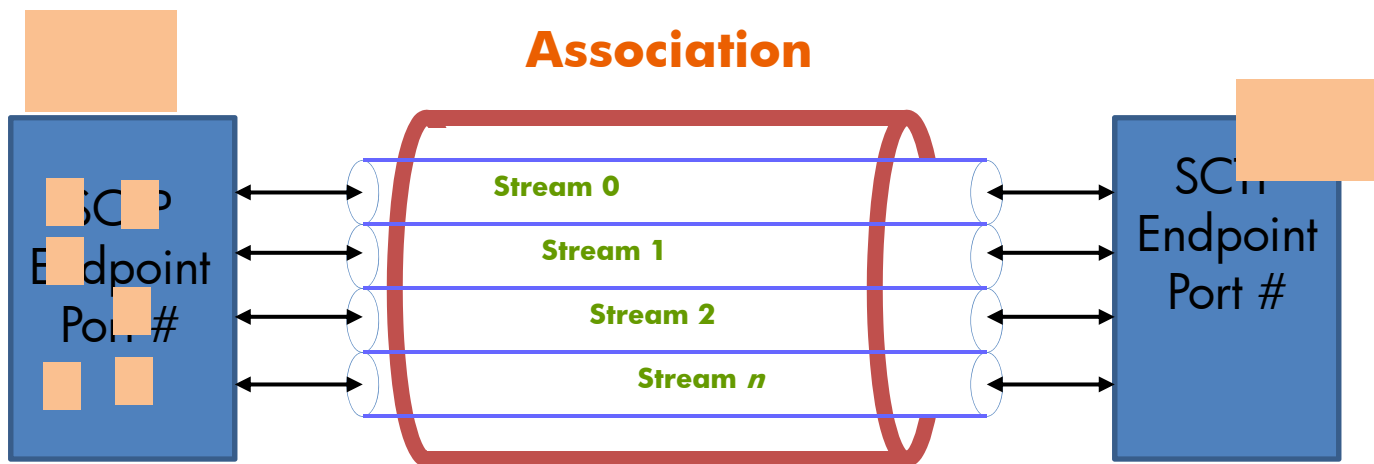
Service/Features	SCTP	TCP	UDP
Message-Oriented	yes	no	yes
Byte-Oriented	no	yes	no
Connection-Oriented	yes	yes	no
Full Duplex	yes	yes	yes
Reliable data transfer	yes	yes	no
Partially-Reliable data transfer	opt	no	no
Ordered data delivery	yes	yes	no
Unordered delivery	yes	no	yes
Flow control	yes	yes	no
Congestion Control	yes	yes	no
ECN Capable	yes	yes	no
Selective Acknowledgments	yes	opt	no
Path MTU discovery	yes	yes	no
Application PDU fragmentation	yes	yes	no
Application PDU bundling	yes	yes	no
Multistreaming	yes	no	no
Multihoming	yes	no	no
Dynamic Multihoming	opt	no	no
SYN flooding attack prevention	yes	no	n/a
Allows half-closed state	no	yes	n/a
Reach-ability check	yes	opt	no
Pseudo-header for checksum	no	yes	yes
Time wait state	no	yes	n/a
Authentication	opt	opt	no
CRC based checksum	yes	no	no



SCTP

Associations and Streaming

- Transport level (RFC 2960) protocol like TCP and UDP
- Streams created within a single SCTP association
- Each stream delivers data independently
- Avoids head-of-line blocking



SCTP

OpenVMS implementation

Based on
Free-BSD

- Well written code and mature code base
- Compiled with
TCPIP\$INTERNET_SERVICES



ENHANCED FTP



Enhanced FTP

Configurations

ANONYMOUS FTP LIGHT

- Enabled via user-based logical TCPIP\$FTP_ANONYMOUS_LIGHT
- Restricts the user's FTP access to directories defined by TCPIP\$FTP_ANONYMOUS_DIRECTORY

FTP Over SSL

- Compliant with RFC 4217
- FTPS is about 2.5 times faster than SFTP



Line Printer Driver (LPD)



Line Printer Driver

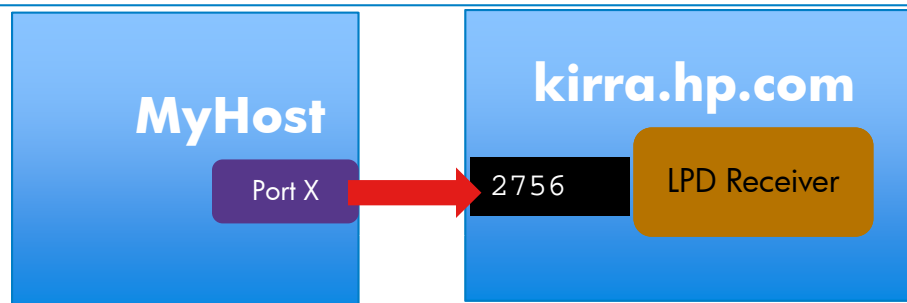
Configurable Port

LPD default port is 515 (RFC 1179)

TCPIP\$PRINTCAP.DAT configuration file

- e.g., Assume LPD listens on remote host:
port = "KIRRA:2756".
Then on the local host, add an entry to TCPIP\$PRINTCAP.DAT

```
PRINTER1: \  
:rm=kirra.hp.com: \  
:rt=2756: \ <etc>
```



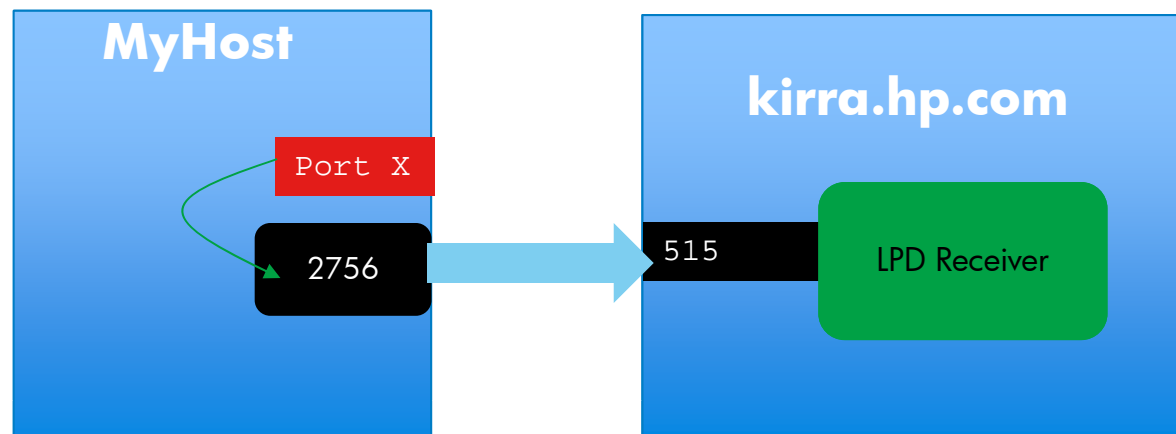
Line Printer Driver

Secure Printing

In TCPIP\$PRINTCAP.DAT set ":rm=localhost:"

- Forward port 2756 over an SSH tunnel to remote port 515, or wherever the remote print server is listening

```
MYHOST$ SSH -"L"2756:localhost:515 kirra.hp.com
```



Mail Enhancements



SMTP

Changes in TCP/IP v5.7

Cluster-aware SMTP

- Load balancing and high availability
- SMTP files/folders - disk visible to all nodes in the cluster.

TCPIP\$SMTP.CONF

- Configuration based on logical names is obsolete
- Roll-over tool ; TCPIP\$SMTP_V57_ROLLOVER.EXE

SMTP Persistent Receiver



Changes in TCP/IP v5.7

TCPIP\$PEERNAME Utility

Displays end-point information

- e.g.,
\$ peername bg1630
Local address: 10.11.12.13, port: 22
Remote address: 10.11.22.210, port: 49573

```
$show sym TCPIP$peername*  
TCPIP$PEERNAME_LOCAL_ADDRESS = "10.11.12.13"  
TCPIP$PEERNAME_LOCAL_PORT    = "22"  
TCPIP$PEERNAME_REMOTE_ADDRESS = "10.11.22.2"  
TCPIP$PEERNAME_REMOTE_PORT    = "49573"
```



NFS Updates



NFS V3 Client on TCP/IP

•NFS V3

Compliant with RFC 1813

Support for larger file sizes (>4GB)

Larger data transfer size (32 KB)

Asynchronous writes to the NFS V3 server using UNSTABLE writes followed by COMMIT requests

Released with TCP/IP 5.7 ECO 2 kit



CIFS Updates



CIFS V1.2+PS001 Enhancements

Improved installation and Automated Configuration

Improved Performance 

ODS 2 Support

Improved File security 

Interoperate with Windows 7 and Windows 2008 R2



CIFS enhancements

Performance enhancements

- Storing file size in an ACE/ACL for non stream format files
- Utility to automatically update file length hint values for Sequential VAR and VFC files
- Open File Caching



CIFS enhancements

File security improvements

- Simplify the use of mask and mode SMB.CONF parameters
- Support for DELETE protection bit with mask and mode SMB.CONF parameters
- Ability to set DELETE protection bit from Windows for OWNER/GROUP/WORLD RMS protection mask



LAN Updates



LAN Update


10Gig

LAN on Motherboard (LOM)

- BL860c i2/BL870c i2/BL890c i2
- Broadcom 57711e
- SYS\$EW57711.EXE driver
- Support for Flex-10 (Virtual Connect Flex-10 10Gb Ethernet Module for c-Class BladeSystem (455880-B21) and DCC

Standup Cards

- HP NC532m PCIe 2-Port 10GbE NIC
- Flex-10 Capable




LAN Update

Flex10

Flex-10 technology is a hardware-based solution that enables users to partition a 10 gigabit Ethernet (10GbE) connection.

Regulate data speed for each partition

Flex-10 technology lets you configure a single 10 Gb Ethernet port to represent four FlexNICs




Flex-10 LAN-on-Motherboard

- Each BL860c i2 motherboard integrates two Broadcom 57711E dual port 10GbE controllers
 - Total of four physical 10GbE ports per motherboard
- With a Flex-10 VC module, the resulting built-in FlexNIC counts are:
 - 16 FlexNICs in a BL860c i2 server
 - 32 FlexNICs in a BL870c i2 server
 - 64 FlexNICs in a BL890c i2 server

Need more?

Add Flex-10 capable mezz cards (like the NC532m)



LAN Update

1Gig

Mezz cards

- HP NC364m PCIe 4-Port 1GbE (Intel)
- HP NC360m PCIe 2-Port 1GbE (Intel)
- BL860c i2/BL870c i2/BL890c i2
- SYS\$EI1000.EXE

HP Integrity rx2800 i2 server

- Core i/o card (LOM)
- AD221A, AD222A, AD393A
- AD337A, AD338A, AD339A



DECnet Plus



DECnet Plus

Enhancements

DECnet/IP communication through SSH – On V8.4

- Uses port forwarding feature of SSH
- Enabled by DECNET_IP_PORT_FORWARD logical

Latest ECO's

ECO's

V8.3
DPLUS ECO 03

V8.3-1H1
DPLUS ECO 03

V8.4
DPLUS ECO 01





Questions / Comments

Office of Customer Programs

OpenVMS.Programs@hp.com

Business Manger

Rohini Madhavan

rohini.madhavan@hp.com



Thank You

