



# Guidelines to Configure Cluster over IP

Table of Contents	
Cluster over IP (IP Cluster Interconnect) .....	2
Design.....	2
PEDRIVER Support for UDP .....	3
TCP/IP Services Boot Time Loading and Initialization .....	3
Availability Manager Support .....	3
Availability .....	3
Configuring Cluster over IP .....	3
Configuration Files.....	4
System Characteristics.....	4
System Requirements.....	4
Satellite Booting Using IP .....	5
Scenario 1: Creating and Configuring a Two-Node Disaster Tolerant Cluster Using IP .....	5
Scenario 2: Adding a new node to a Cluster over IP.....	15
Scenario 3: Adding a new Node to a Cluster over IP with a Shared System Disk .....	17
Scenario 4: Adding an Integrity server Satellite Node to a Cluster over IP .....	24
Scenario 5: Adding an Integrity server Node to a Cluster over IP with Logical LAN Failover set .....	33
Appendix A - Cluster over IP Commands .....	37
Using SDA to Monitor Cluster Communications.....	43
Appendix B Restrictions and Known Problems in this Release .....	47
For more information.....	49

# Cluster over IP (IP Cluster Interconnect)

OpenVMS Version 8.4 has been enhanced with the Cluster over IP (Internet Protocol) feature. Cluster over IP provides the ability to form clusters beyond a single LAN or VLAN segment using industry standard Internet Protocol. This feature provides improved disaster tolerant capability.

System managers also have the ability to manage or monitor OpenVMS cluster that uses IP for cluster communication using SCACP management utility. Cluster protocol (SCS also known as SCA) over LAN is provided by Port emulator driver (PEDRIVER). PEDRIVER uses User datagram protocol (UDP) and IP in addition to directly using 802.3 interfacing with LAN for cluster communication as shown in Figure 1-0. The datagram characteristics of UDP combined with PEDRIVER's inbuilt reliable delivery mechanism provide a standard cluster communication over which SYSAP (system level application) communicate between two cluster nodes.

Cluster over IP is an optional feature that can be enabled in addition to the traditional LAN based communication. However, if both LAN and IP mode of communication exist in a cluster, PEDRIVER is designed to use the LAN interconnect over IP for cluster communication.

**Note:** OpenVMS cluster over IP and IP cluster Interconnect (IPCI) terms are interchangeably used in the document and refers to using TCP/IP stack for cluster communication.

---

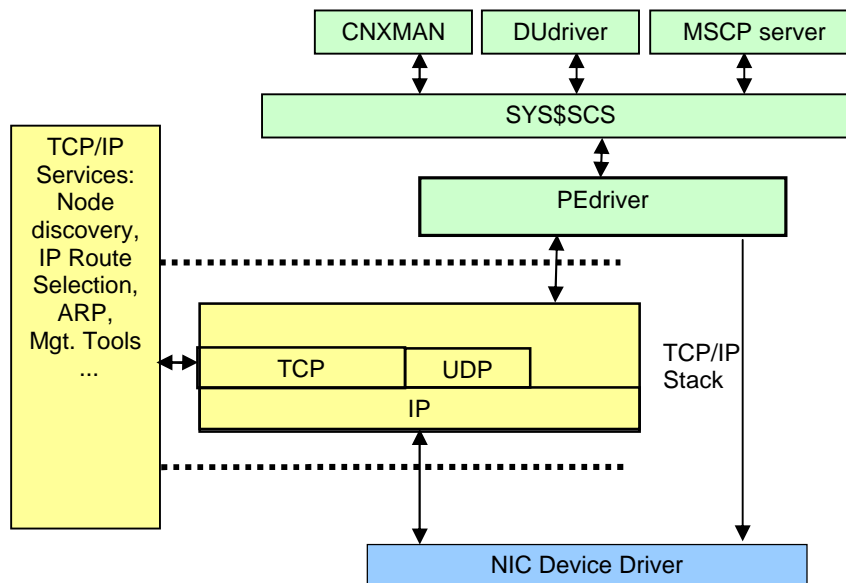
## Design

Cluster over IP solution is an integration of the following:

- PEDRIVER support for UDP protocol
- TCP/IP Services boot time loading and initialization

Figure 1.0 shows the cluster communication using IP.

Figure 1.0 Cluster Communication Design Using IP



## PEDRIVER Support for UDP

PEDRIVER is extended to support IP with the UDP protocol, which has the same packet delivery characteristics as that of 802 LANs. PEDRIVER implements the transport layer of NISCA, which has inbuilt delay probing, retransmission, reliable delivery for sequenced messages, implements datagram service. The variable buffer size for block transfers for I/O is suitable for cluster traffic.

PEDRIVER uses 802 multicast for discovering LAN connected nodes, and for the keep-alive mechanism. IP multicast maps 1:1 onto the existing LAN discovery. It is used as the preferred node discovery mechanism. Administratively scoped IP multicast address can be used for exchanging hello packets. IP unicast address of nodes can also be used for node discovery and to join the cluster in environments where IP multicast can not be used.

The SCACP Management interface has been added with the ability to control, configure, and monitor PEDRIVER use of IP interfaces for cluster communication.

## TCP/IP Services Boot Time Loading and Initialization

To ensure that cluster communication is available in an IP only network environment, it is essential to have TCP/IP stack loaded when the cluster formation starts. This also retains the existing functionality of cluster formation of OpenVMS clusters. Normal booting sequence includes loading of LAN drivers followed by PEDRIVER. TCP/IP drivers are loaded when TCP/IP services are started. If cluster over IP is enabled, LAN, TCP/IP exelets, and PEDRIVER are loaded sequentially. Once the system comes up, TCP/IP services can be started to use other TCP/IP components, such as TELNET, FTP and so on.

## Availability Manager Support

Cluster over IP provides the ability to form clusters beyond a single LAN or VLAN segment using industry standard Internet protocol. Cluster over IP provides improved disaster tolerant capability. Availability Manager Version 3.1 has been enhanced with the ability to support OpenVMS cluster over IP functionality. This new release of Availability Manager provides functionality to manage and monitor LAN or IP path (channels) data; IP interface (IP bus) used for cluster communication and virtual circuit formed using IP channels.

## Availability

The ability to create a logical LAN failover set using IP for cluster communication provides high availability systems. The nodes will be able to resume if a local LAN card fails, as it will switchover to another interface configured in the logical LAN failover set. For a complete description of creating a logical LAN failover set, see *Guidelines for OpenVMS Cluster Configurations*.

## Configuring Cluster over IP

IP multicast can be used for node discovery within the same IP multicast domain. The IP multicast address is administratively scoped. The remote nodes not in the IP multicast domain can use the IP unicast technique to join the cluster and send hello packets. The NISCS\_USE\_UDP SYSGEN parameter enables IP as interconnect and can be modified while executing the CLUSTER\_CONFIG\_LAN.COM.

## Configuration Files

SYS\$SYSTEM:PE\$IP\_CONFIG.DAT and SYS\$SYSTEM:TCPIP\$CLUSTER.DAT are the two configuration files used for configuring Cluster with IP interconnect. These files are loaded during the boot process and provide the necessary configuration details for Cluster over IP. Both these files are generated when a node is configured to be a member of the cluster and if cluster over IP is enabled during the configuration.

SYS\$SYSTEM:PE\$IP\_CONFIG.DAT includes the optional IP multicast and IP unicast addresses of the nodes of the cluster. IP multicast messages are used for discovering a node within the same IP multicast domain. Remote nodes in a different IP multicast domain can use the IP unicast messaging technique to join the cluster. SYS\$SYSTEM:PE\$IP\_CONFIG.DAT can be common for all the nodes of a cluster.

SYS\$SYSTEM:TCPIP\$CLUSTER.DAT contains the IP interface name and IP addresses on which cluster communication is enabled. It also includes the TCP/IP route information.

SYS\$SYSTEM:TCPIP\$CLUSTER.DAT is unique for each node in a cluster.

## System Characteristics

The existing functionalities of OpenVMS Clusters continue to exist with Cluster over IP.

It has the following characteristics:

- Cluster over IP does not require any new hardware to use TCP/IP stack as interconnect.
- UDP protocol is used for cluster communication.
- The PEDRIVER includes delay probing technique that helps reduce latency in the IP network by selecting a path with the least latency.
- The OpenVMS Cluster feature of rolling upgrades to the new version without a cluster reboot is retained.
- Provides interoperability with servers running earlier versions of OpenVMS Clusters that are LAN based. Cluster over IP is available only with OpenVMS Version 8.4. Hence, if the node requires IP interconnect to be part of the cluster, then all the nodes of the cluster must be running OpenVMS Version 8.4 and HP TCP/IP Services for OpenVMS, Version 5.7.
- At the boot time, LAN, TCP/IP, PEDRIVER are started sequentially.
- PEDRIVER automatically detects and creates an IP channel for communication between two nodes.
- Cluster over IP feature can be optionally enabled by running the CLUSTER\_CONFIG\_LAN.COM.
- IP address used for cluster communication must be primary static address of the interface.

## System Requirements

The following software is required to support Cluster over IP interconnect:

- OpenVMS Version 8.4 for Integrity servers or OpenVMS Alpha Version 8.4
- HP TCP/IP services for OpenVMS Version 5.7
- Static IP addresses, IP unicast and optional IP Multicast for node discovery

**Note:** Ensure that the TCP/IP software is configured before configuring cluster over IP. To ensure that network and TCP/IP is configured properly, use the PING utility and ping the node from outside the subnet.

---

## Satellite Booting Using IP

Satellite booting is supported on both Integrity servers and Alpha systems. The Integrity server satellite node must be in the same LAN on which the boot server resides. The Alpha satellite node must be in the same LAN as its disk server.

The Alpha console uses the MOP protocol for network load of satellite systems. Because the MOP protocol is non-routable, the satellite boot server or servers and all satellites booting from them must reside in the same LAN. In addition, the boot server must have at least one LAN device enabled for cluster communications to permit the Alpha satellite nodes to access the system disk.


**See Appendix** "Cluster over IP Commands" for the commands used to configure and monitor cluster systems interconnected with IP.

---

## Scenario 1: Creating and Configuring a Two-Node Disaster Tolerant Cluster Using IP

Cluster over IP can be used to create and configure a two-node disaster tolerant (DT) cluster. Node ORCHID is a standalone node at SITE A and node TULIP at SITE B, is already a member (the only member) of the cluster. In this scenario, Cluster over IP is not configured in TULIP. SITE A and SITE B can be in the same or different LAN, building or any other geographical location. It is required to have IP connectivity between SITE A and SITE B and must be within the supported intersite distance.

Figure 2.0 Configuring a Two-Node Disaster Tolerant Cluster



### Step 1. Configuring Node TULIP to Enable Cluster over IP

To configure the node TULIP (OpenVMS Alpha node) for enabling the Cluster over IP feature, execute the CLUSTER\_CONFIG\_LAN.COM procedure on node TULIP and select the appropriate option as illustrated:

```
TULIP$ @SYS$MANAGER:CLUSTER_CONFIG_LAN
```

```
Cluster Configuration Procedure  
CLUSTER_CONFIG_LAN Version V2.79  
Executing on an Alpha System
```

DECnet-Plus is installed on this node.  
Alpha satellites will use LANCP, not DECnet, for MOP downline loading.

Enter a "?" for help at any prompt. If you are familiar with the execution of this procedure, you may want to mute extra notes and explanations by invoking it with "@CLUSTER\_CONFIG\_LAN BRIEF".

TULIP is an Alpha system and currently a member of a cluster so the following functions can be performed:

MAIN Menu

1. ADD an Alpha node to the cluster.
2. REMOVE a node from the cluster.
3. CHANGE a cluster member's characteristics.
4. CREATE a duplicate system disk for TULIP.
5. MAKE a directory structure for a new root on a system disk.
6. DELETE a root from a system disk.
7. EXIT from this procedure.

Enter choice [7]: 3 .....[1]  
CHANGE Menu

1. Enable TULIP as a boot server.
2. Disable TULIP as a boot server.
3. Enable a quorum disk for TULIP
4. Disable a quorum disk for TULIP.
5. Enable TULIP as a disk server.
6. Disable TULIP as a disk server.
7. Change TULIP's ALLOCLASS value.
8. Enable TULIP as a tape server.
9. Disable TULIP as a tape server.
10. Change TULIP's TAPE\_ALLOCLASS value.
11. Change an Alpha satellite node's LAN adapter hardware address.
12. Enable Cluster Communication using IP on TULIP.
13. Disable Cluster Communication using IP on TULIP.
14. Enable the LAN for cluster communications on TULIP.
15. Disable the LAN for cluster communications on TULIP.
16. Enable Memory Channel for cluster communications on TULIP.
17. Disable Memory Channel for cluster communications on TULIP.
18. Change TULIP's shared SCSI port allocation class value.
19. Return to MAIN menu.

Enter choice [19]: 12.....[2]

ENABLE IP for cluster communications (Y/N)? y.....[3]  
UDP port number to be used for Cluster Communication over IP[49152]?.....[4]  
Enable IP multicast for cluster communication(Y/N)[Y]? Y.....[5]  
What is the IP multicast address[239.242.7.193]?.....[6]  
What is the TTL (time to live) value for IP multicast packets [32] ?.....[7]  
Do you want to enter unicast address(es)(Y/N)[Y]?.....[8]  
What is the unicast address[Press [RETURN] to end the list]? 10.0.1.2 .[9]  
What is the unicast address[Press [RETURN] to end the list]?.....[10]

\*\*\*\*\*  
Cluster Communications over IP has been enabled. Now

CLUSTER\_CONFIG\_LAN will run the SYS\$MANAGER:TCPIP\$CONFIG procedure. Please select the IP interfaces to be used for Cluster Communications over IP (IPCI). This can be done selecting "Core Environment" option from the main menu followed by the "Interfaces" option. You may also use this opportunity to configure other aspects.  
\*\*\*\*\*

Press Return to continue .

Checking TCP/IP Services for OpenVMS configuration database files.

HP TCP/IP Services for OpenVMS Configuration Menu

Configuration options:

- 1 - Core environment
- 2 - Client components
- 3 - Server components
- 4 - Optional components
- 5 - Shutdown HP TCP/IP Services for OpenVMS
- 6 - Startup HP TCP/IP Services for OpenVMS
- 7 - Run tests
- A - Configure options 1 - 4
- [E] - Exit configuration procedure

Enter configuration option: 1 .....[11]

HP TCP/IP Services for OpenVMS Core Environment Configuration Menu

Configuration options:

- 1 - Domain
- 2 - Interfaces
- 3 - Routing
- 4 - BIND Resolver
- 5 - Time Zone
- A - Configure options 1 - 5
- [E] - Exit menu

Enter configuration option: 2.....[12]

HP TCP/IP Services for OpenVMS Interface & Address Configuration Menu

Hostname Details: Configured=TULIP, Active=TULIP

Configuration options:

- 0 - Set The Target Node (Current Node: TULIP)
- 1 - IE0 Menu (EIA0: TwistedPair 100mbps)
- 2 - 10.0.2.2/23 TULIP Configured,Active
- 3 - IE1 Menu (EIB0: TwistedPair 100mbps)
- 4 - 10.0.2.224/23 \*noname\* Configured,Active
- I - Information about your configuration
- [E] - Exit menu

Enter configuration option: 2 .....[13]

HP TCP/IP Services for OpenVMS Address Configuration Menu (Node: TULIP)

IE0 10.0.2.2/23 TULIP Configured, Active IE0

Configuration options

- 1 - Change address
- 2 - Set "TULIP" as the default hostname
- 3 - Delete from configuration database
- 4 - Add to IPCI database
- 5 - Deactivate from live system
- 6 - Add standby aliases to configuration database (for failSAFE IP)
- [E] - Exit menu

Enter configuration option: 4.....[14]

Updated Interface in IPCI configuration file:

SYS\$SYSROOT:[SYSEXE]TCPIP\$CLUSTER.DAT;

Updated Default Route in IPCI configuration file:

SYS\$SYSROOT:[SYSEXE]TCPIP\$CLUSTER.DAT;

Added address IE1:10.0.2.2 to IPCI database

HP TCP/IP Services for OpenVMS Interface & Address Configuration Menu

Hostname Details: Configured=TULIP, Active=TULIP

Configuration options:

- 0 - Set The Target Node (Current Node:TULIP)
- 1 - IE0 Menu (EIA0: TwistedPair 100mbps)
- 2 - 10.0.2.2/23 TULIP Configured,IPCI,Active
- 3 - IE1 Menu (EIB0: TwistedPair 100mbps)
- 4 - 10.0.2.224/23 \*noname\* Configured,Active
- I - Information about your configuration
- [E]- Exit menu

Enter configuration option: E.....[15]

HP TCP/IP Services for OpenVMS Core Environment Configuration Menu

Configuration options:

- 1 - Domain
- 2 - Interfaces
- 3 - Routing
- 4 - BIND Resolver
- 5 - Time Zone
- A - Configure options 1 - 5
- [E] - Exit menu

Enter configuration option: E

HP TCP/IP Services for OpenVMS Configuration Menu

Configuration options:

- 2 - Client components
- 3 - Server components
- 4 - Optional components
- 5 - Shutdown HP TCP/IP Services for OpenVMS
- 6 - Startup HP TCP/IP Services for OpenVMS
- 7 - Run tests
- A - Configure options 1 - 4
- [E] - Exit configuration procedure

Enter configuration option: E

The configuration procedure has completed successfully.

TULIP has been enabled for IP communications.....[16]  
Please run AUTOGEN to reboot TULIP:

TULIP\$ @SYS\$UPDATE:AUTOGEN GETDATA REBOOT .....[17]

Field	Description
[1]	TULIP is a single-member cluster without Cluster over IP enabled. The cluster member characteristics can be changed to enable cluster over IP for this node by selecting option 3.
[2]	Select option 12 to enable cluster over IP. By selecting this option, the SYSGEN parameter, NISCS_USE_UDP is set to 1, which enables the PEDRIVER to use IP for cluster communication. This requires reboot of the node. If LAN is not already selected as the cluster interconnect, this option sets NISCS_LOAD_PEA0 to 1 to load PEDRIVER during the next reboot.
[3]	Enable IP for cluster communication.
[4]	The UDP port number used for cluster communication. The UDP port number must be same on all members of the cluster. Also, ensure that there is no other cluster in your environment using the same UDP port number and this port number must not be used by any other application.
[5]	You can enable IP multicast for cluster communication if your environment allows IP multicast traffic between cluster nodes. Check with your network administrator, if IP multicasting is enabled in your environment.
[6]	Enter the IP multicast address for the cluster, if IP multicasting is enabled. By default, the IP multicast address is selected from the administratively scoped IP multicast address ranging from 239.242.x.y. The last two octets x and y are generated based on the cluster group number. In the above example cluster group number is 1985 and can be calculated as follows: X= 1985/256 Y= 1985 - (256 *x) The system administrator can override the default multicast address by a unique address for their environment.
[7]	TTL is the time to live for IP multicast packets. It specifies the number of hops allowed for IP multicast packets.
[8]	Enter "Yes" to enter the IP unicast addresses for the remote nodes of the cluster, which are not reachable using IP multicast address.

[9]	<p>In this example, 10.0.1.2 is the IP unicast address for the node ORCHID. Although, the IP multicast is selected, ORCHID's IP address is entered because the IP multicast connectivity between SITE A and SITE B is presumed to be not existing in this example.</p> <p>Note: Enter the list of IP address of the cluster. All the information entered in [4], [6], [7], and [9] are entered into the SYS\$SYSTEM:PE\$IP_CONFIG.DAT file.</p> <p>The PE\$IP_CONFIG.DAT file is generated as shown in the following example.</p> <p>Also, to allow the remote node to join the cluster, the Unicast list in the PE\$IP_CONFIG.DAT on the local node must contain the IP address of the remote node.</p> <p>In this example, TULIP must have ORCHID's IP address and ORCHID must have TULIP's IP address. SYSTEM:PE\$IP_CONFIG.DAT on node TULIP is as shown:</p> <pre> ! CLUSTER_CONFIG_LAN creating for CHANGE operation on 10-JUL- 2008 14:14:06.57 multicast_address=239.242.7.193 ttl=32 udp_port=49152 unicast=10.0.1.2 </pre>
[10]	Press Return after entering the unicast list.
[11]	<p>CLUSTER_CONFIG_LAN.COM invokes TCPIP\$CONFIG.COM to configure the IP interfaces used for cluster communication. Select the core environment option.</p> <p>Assuming that TCP/IP is already configured, the node can be pinged from outside the subnet.</p>
[12]	Select the interface option from the core environment menu.
[13]	Select the appropriate interface for cluster communication. In this example, option 2 is selected.
[14]	<p>In the configuration option for the selected address, select option 4 to add to the Cluster over IP database. The interface information along with the default route is entered in the TCPIP\$CLUSTER.DAT as shown:</p> <pre> interface=IE0,EIA0,10.0.2.2,255.255.254.0 default_route=10.0.2.1 </pre>
[15]	Exit from the TCP/IP configuration procedure, which returns to CLUSTER_CONFIG_LAN.COM.
[16]	Proceed with cluster configuration.
[17]	After rebooting the system, run AUTOGEN. PEDRIVER in ORCHID will start using IP in addition to LAN for cluster communication and must be able to join TULIP.

## Step 2. Configuring Node ORCHID as a Cluster Member

To configure ORCHID with Cluster over IP enabled, execute CLUSTER\_CONFIG\_LAN.COM on node ORCHID and select the appropriate option as shown:

```
ORCHID$ @SYS$MANAGER:CLUSTER_CONFIG_LAN
```

```

Cluster Configuration Procedure
CLUSTER_CONFIG_LAN Version V2.79
Executing on an IA64 System

```

```

DECnet-Plus is installed on this node.
IA64 satellites will use TCP/IP BOOTP and TFTP services for downline
loading.

```

```
TCP/IP is installed and running on this node.
```

```

Enter a "?" for help at any prompt. If you are familiar with the
execution of this procedure, you may want to mute extra notes and
explanations by invoking it with "@CLUSTER_CONFIG_LAN BRIEF".

```

```
This IA64 node is not currently a cluster member.
```

```
MAIN Menu
```

1. ADD ORCHID to existing cluster, or form a new cluster.
2. MAKE a directory structure for a new root on a system disk.
3. DELETE a root from a system disk.
4. EXIT from this procedure.

Enter choice [4]: 1 .....[1]

Is the node to be a clustered node with a shared SCSI/FIBRE-CHANNEL bus (Y/N)? N

What is the node's SCS node name? ORCHID

IA64 node, using LAN/IP for cluster communications. PEDRIVER will be loaded.

No other cluster interconnects are supported for IA64 nodes.

Enter this cluster's group number: 1985

Enter this cluster's password:

Re-enter this cluster's password for verification:

ENABLE IP for cluster communications (Y/N)? Y .....[2]

UDP port number to be used for Cluster Communication over

IP[49152]?.....[3]

Enable IP multicast for cluster communication(Y/N)[Y]? Y .....[4]

What is IP the multicast address[239.242.7.193]? .....[5]

What is the TTL (time to live) value for IP multicast packets [32] ?.....[6]

Do you want to enter unicast address(es)(Y/N)[Y]?.....[7]

What is the unicast address[Press [RETURN] to end the list]?

10.0.2.2.....[8]

What is the unicast address[Press [RETURN] to end the list]? .....[9]

```

*****
Cluster Communications over IP has been enabled. Now
CLUSTER_CONFIG_LAN will run the SYS$MANAGER:TCPIP$CONFIG procedure.
Please select the IP interfaces to be used for Cluster
Communications over IP (IPCI). This can be done selecting "Core
Environment" option from the main menu followed by the "Interfaces"
option. You may also use this opportunity to configure other
aspects.
*****

```

Press Return to continue ...

#### TCP/IP Network Configuration Procedure

This procedure helps you define the parameters required to run HP TCP/IP Services for OpenVMS on this system.

%TCPIP-I-IPCI, TCP/IP Configuration is limited to IPCI.

-TCPIP-I-IPCI, Rerun TCPIP\$CONFIG after joining the cluster.

#### HP TCP/IP Services for OpenVMS Interface & Address Configuration Menu

Hostname Details: Configured=Not Configured, Active=nodeg

Configuration options:

- 0 - Set The Target Node (Current Node: ORCHID)
- 1 - IE0 Menu (EIA0: TwistedPair 100mbps)

- 2 - IE1 Menu (EIB0: TwistedPair 100mbps)
- [E] - Exit menu

Enter configuration option: 1 .....[10]

\* IPCI Address Configuration \*

Only IPCI addresses can be configured in the current environment.  
 After configuring your IPCI address(es) it will be necessary to  
 run TCPIP\$CONFIG once your node has joined the cluster.

IPv4 Address may be entered with CIDR bits suffix.  
 E.g. For a 16-bit netmask enter 10.0.1.1/16

Enter IPv4 Address []:10.0.1.2/24 .....[11]  
 Default netmask calculated from class of IP address: 255.0.0.0

IPv4 Netmask may be entered in dotted decimal notation,  
 (e.g. 255.255.0.0), or as number of CIDR bits (e.g. 16)

Enter Netmask or CIDR bits [255.0.0.0]: 255.255.254.0.....[12]

Requested configuration:

```

Node      : ORCHID
Interface: IE0
IPCI      : Yes
Address   : 10.0.1.2/23
Netmask   : 255.255.254.0 (CIDR bits: 23)
  
```

\* Is this correct [YES]:

Updated Interface in IPCI configuration file:

SYS\$SYSROOT:[SYSEXE]TCPIP\$CLUSTER.DAT; .....[13]

HP TCP/IP Services for OpenVMS Interface & Address Configuration Menu

Hostname Details: Configured=Not Configured, Active=ORCHID

Configuration options:

- 0 - Set The Target Node (Current Node: ORCHID)
- 1 - IE0 Menu (EIA0: TwistedPair 100mbps)
- 2 - 10.0.1.2 /23 \*noname\* IPCI
- 3 - IE1 Menu (EIB0: TwistedPair 100mbps)
- [E] - Exit menu

Enter configuration option: E.....[14]

Enter your Default Gateway address []: 10.0.1.1.....[15]

\* The default gateway will be: 10.0.1.1. Correct [NO]: YES

Updated Default Route in IPCI configuration file:

SYS\$SYSROOT:[SYSEXE]TCPIP\$CLUSTER.DAT;

TCPIP-I-IPCIDONE, Finished configuring IPCI address information.....[16]

Will ORCHID be a boot server [Y]? .....[17]

TCP/IP BOOTP and TFTP services must be enabled on IA64 boot nodes.

Use SYS\$MANAGER:TCPIP\$CONFIG.COM on ORCHID to enable BOOTP and TFTP services after ORCHID has booted into the cluster.

Enter a value for ORCHID's ALLOCLASS parameter [7]:  
Does this cluster contain a quorum disk [N]?

The EXPECTED\_VOTES system parameter of members of a cluster indicates the total number of votes present when all cluster members are booted, and is used to determine the minimum number of votes (QUORUM) needed for cluster operation.

EXPECTED\_VOTES value for this cluster: 1

Warning: Setting EXPECTED\_VOTES to 1 allows this node to boot without being able to see any other nodes in the cluster. If there is another instance of the cluster in existence that is unreachable via SCS but shares common drives (such as a Fibrechannel fabric) this may result in severe disk corruption.

Do you wish to re-enter the value of EXPECTED\_VOTES [Y]? n

The use of a quorum disk is recommended for small clusters to maintain cluster quorum if cluster availability with only a single cluster node is a requirement.

For complete instructions, check the section on configuring a cluster in the "OpenVMS Cluster Systems" manual.

WARNING: ORCHID will be a voting cluster member. EXPECTED\_VOTES for this and every other cluster member should be adjusted at a convenient time before a reboot. For complete instructions, check the section on configuring a cluster in the "OpenVMS Cluster Systems" manual.

Execute AUTOGEN to compute the SYSGEN parameters for your configuration and reboot ORCHID with the new parameters. This is necessary before ORCHID can become a cluster member.

Do you want to run AUTOGEN now [Y]? N

Please run AUTOGEN to reboot ORCHID:

ORCHID\$ @SYS\$UPDATE:AUTOGEN GETDATA REBOOT .....[18]

Field	Description
[1]	Node ORCHID is currently a standalone, Integrity server node and is made as a member of a cluster. Only LAN or IP is used for cluster communication and no other interconnect is supported.
[2]	Select IP for cluster communication in addition to LAN by entering "YES". The SYSGEN parameter, NISCS_USE_UDP is set to 1 and PEDRIVER uses IP in addition to LAN for cluster communication when the node is rebooted.
[3]	The UDP port number to be used for cluster communication. The UDP port number must be same on all members of the cluster. Also, ensure that there is no other cluster in your environment using the same UDP port number and this port number must not be used by any other application.

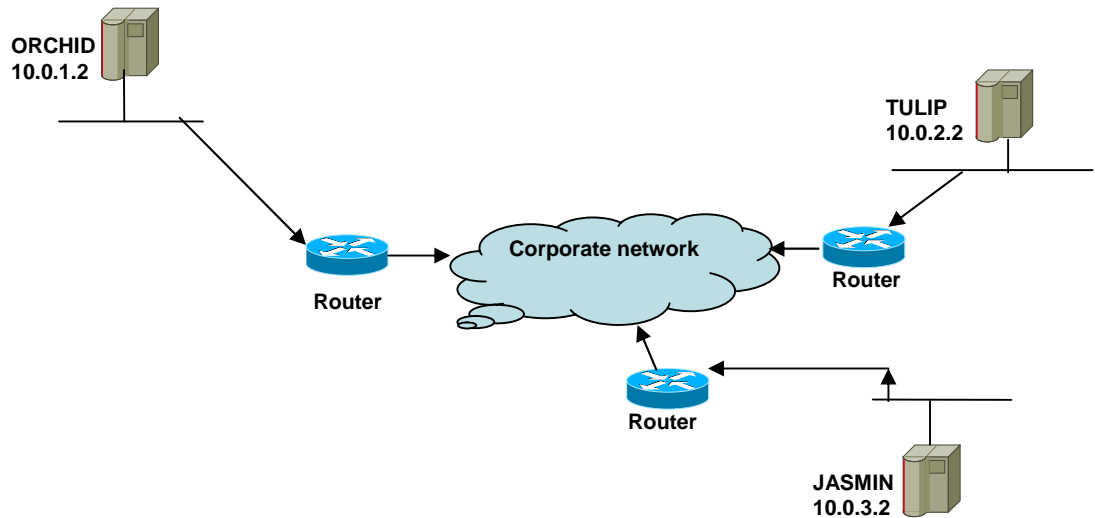
[4]	You can enable IP multicast for cluster communication if your environment allows IP multicast traffic between cluster nodes. Check with your network administrator to see if IP multicasting is enabled in your environment.
[5]	Enter the IP multicast address for cluster, if IP multicasting is enabled. By default, the IP multicast address is selected from the administratively scoped IP multicast address range of 239.242.x.y. The last two octets x and y are generated based on the cluster group number. In the above example cluster group number is 1985 and is calculates as follows: $X = 1985 / 256$ $Y = 1985 - (256 * x)$ System Administrator can override the default multicast address by a unique address for their environment.
[6]	TTL is the time to live for IP multicast packets. It specifies the number of hops allowed for IP multicast packets.
[7]	Enter "Yes" to enter the IP unicast address for remote nodes of the cluster which are not reachable using IP multicast address.
[8]	In this example, 10.0.2.2 is the IP unicast address of node TULIP. Although, IP multicast is selected, TULIP's IP address is entered because the IP multicast connectivity between SITE A and SITE B is presumed to be non-existent in this example. NOTE: Enter the list of IP unicast address of the cluster. All the information entered in [2], [3], [5], [6], and [7] are entered in the SYS\$SYSTEM:PE\$IP_CONFIG.DAT file. The PE\$IP_CONFIG.DAT file is generated as shown in following example. Also, the Unicast list in PE\$IP_CONFIG.DAT in the local node must contain the remote node IP address for the local node to allow the remote node to join the cluster. In this example, ORCHID must have TULIP's IP address and TULIP must have ORCHID's IP address. SYSTEM:PE\$IP_CONFIG.DAT in node ORCHID <pre>! CLUSTER_CONFIG_LAN creating for CHANGE operation on 10-JUL-2008 14:14:06.57 multicast_address=239.242.7.193 ttl=32 udp_port=49152 unicast=10.0.2.2</pre>
[9]	Press Return after entering the Unicast list.
[10]	CLUSTER_CONFIG_LAN.COM invokes TCPIP\$CONFIG.COM to configure the IP interfaces used for cluster communication. Currently, ORCHID is a standalone node, when TCPIP\$CONFIG is invoked by the CLUSTER_CONFIG_LAN procedure, TCP/IP configuration is limited to IPCI. The interface, IE0 is selected for enabling cluster communications.
	Note: TCPIP\$CONFIG must be invoked after joining the cluster for other TCP/IP configuration such as, FTP, TELNET.
[11]	IPv4 address for the IE0 interface is 10.0.1.2
[12]	Network mask for the IE0 interface is 255.255.254.0
[13]	The IE0 interface information along with network mask is entered in the TCPIP\$CLUSTER.DAT file.
[14]	Exit the interface menu after selecting the interface for cluster communication.
[15]	The default gateway address for the interface IE0 is entered. Only one default gateway address is allowed for IPCI communication.
[16]	After the interface and default gateway are selected, TCPIP\$CONFIG updates the TCPIP\$CLUSTER.DAT with the default route or gateway information. This also completes the TCPIP\$CONFIG required for cluster communications using IP. The interface information along with the default route is entered in the TCPIP\$CLUSTER.DAT file as shown: <pre>interface=IE0,EIA0,10.0.1.2,255.255.254.0 default_route=10.0.1.1</pre>
[17]	Proceed with cluster configuration.

[18] After rebooting the system, run AUTOGEN. PEDRIVER in ORCHID will start using IP in addition to LAN for cluster communication and must be able to join TULIP.

## Scenario 2: Adding a new node to a Cluster over IP

This section describes how to add a new node, JASMIN to an existing two-node cluster. Nodes, ORCHID and TULIP are currently members of a two-node cluster, which are at SITE A and SITE B. For more information about configuring a node with IP as interconnect, see Scenario 1. Node JASMIN is currently a standalone node at SITE C with IP connectivity to both SITE A and SITE B.

Figure 3.0 Adding a new node to Cluster over IP



### Step 1. Ensuring IP Connectivity

Ensure that the IP connectivity between the node JASMIN and the nodes ORCHID and TULIP is working fine. Use the TCP/IP PING utility to test the IP connectivity between JASMIN and other nodes, ORCHID and TULIP.

If PING fails, set up TCP/IP configuration properly so that node JASMIN can ping both ORCHID and TULIP.

### Step 2. Executing the CLUSTER\_CONFIG\_LAN.COM

Execute CLUSTER\_CONFIG\_LAN.COM on node JASMIN. Because, the node JASMIN is a standalone node, complete the procedure described in Step 2 of Scenario 1.

Complete the sequence of steps provided in the following example while entering the Unicast list.

```
Do you want to enter unicast address(es)(Y/N)[Y]?.  
What is the unicast address[Press [RETURN] to end the list]? 10.0.3.2  
What is the unicast address[Press [RETURN] to end the list]? 10.0.2.2
```

What is the unicast address[Press [RETURN] to end the list]? 10.0.1.2  
.....[1]

What is the unicast address[Press [RETURN] to end the list]?

Field	Description
[1]	Enter the IP address of JASMIN ORCHID and TULIP while configuring the node JASMIN.

SYS\$SYSTEM:PE\$IP\_CONFIG.DAT file generated in node JASMIN shown below

```
! CLUSTER_CONFIG_LAN creating for CHANGE operation on 10-JUL-2008
14:14:06.57
multicast_address=239.242.7.193
ttl=32
udp_port=49152
unicast=10.0.3.2
Unicast=10.0.2.2
Unicast=10.0.1.2
```

**Note:** The unicast list must be consistent in all nodes of the cluster. Hence, while entering the unicast list in JASMIN, enter the IP addresses of all the three nodes of the cluster (that is, JASMIN, ORCHID and TULIP). You can also enter the local nodes IP addresses along with the unicast list as it facilitates system management.

---

### Step 3. Completing the Configuration Procedure

Continue to run the CLUSTER\_CONFIG\_LAN.COM to complete the cluster configuration procedure, For more information, see Scenario 1.

### Step 4. Updating the PE\$IP\_CONFIG.DAT File

To ensure that the nodes join the cluster, it is required to have PE\$IP\_CONFIG.DAT consistent through all the members of the cluster. Copy the SYS\$SYSTEM:PE\$IP\_CONFIG.DAT file that is created on node JASMIN to the other nodes, ORCHID, and TULIP.

### Step 5. Refreshing the Unicast List

On both ORCHID and TULIP, to update the new unicast list in the PE\$IP\_CONFIG.DAT file, enter the following command for PEDRIVER:

```
$MC SCACP RELOAD
```

You can also use SYSMAN and run the command clusterwide.

**Note:** The following rule is applicable when IP unicast address is used for node discovery. A node will be allowed to join the cluster only if its IP address is present in the existing members of the SYS\$SYSTEM:PE\$IP\_CONFIG.DAT file.

---

### Step 6. Running AUTOGEN and Rebooting the Node

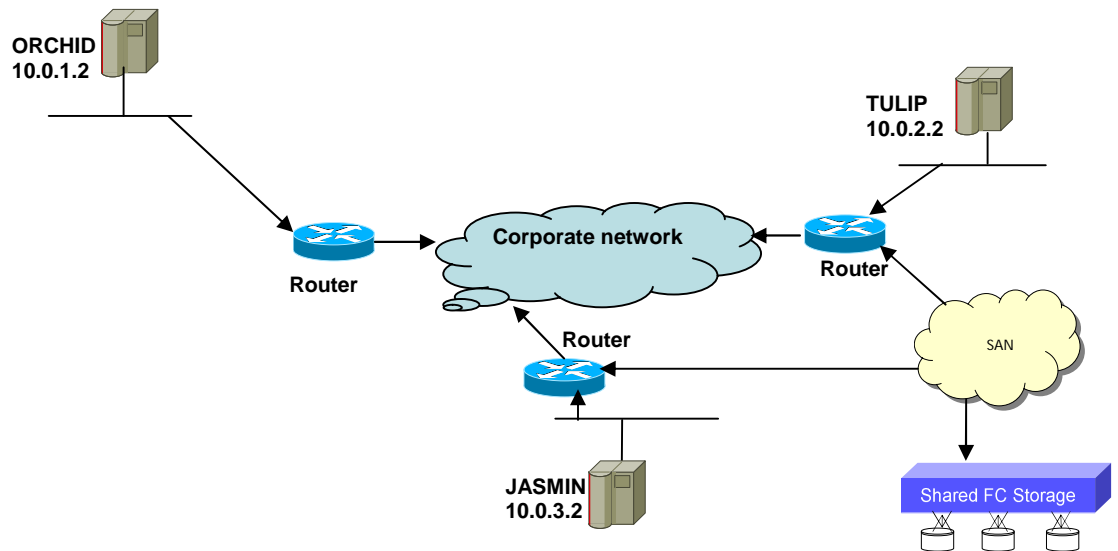
After the first boot of JASMIN, AUTOGEN.COM will run automatically to join the existing cluster consisting of nodes ORCHID and LOTUS.

```
JASMIN$ @SYS$UPDATE:AUTOGEN GETDATA REBOOT
```

## Scenario 3: Adding a new Node to a Cluster over IP with a Shared System Disk

This section describes how to add a new node JASMIN that has a shared system disk of TULIP. ORCHID and TULIP are currently members of two-node cluster which are at SITE A and SITE B.

Figure 4.0 Cluster over IP with a Shared System Disk



### Step 1. Obtaining the Interface Information

Node JASMIN is an OpenVMS Alpha node and is directly connected to the system disk of one of the node TULIP. In this configuration, node JASMIN is connected in network, but not yet booted.

To configure cluster over IP, the interface information of JASMIN is required. This information can be obtained from the '>>>' prompt on JASMIN by executing the following command:

```
P00>>>SHOW DEVICE
dga5245.1003.0.3.0      $1$DGA5245  COMPAQ HSV110 (C)COMPAQ  3028
dga5245.1004.0.3.0      $1$DGA5245  COMPAQ HSV110 (C)COMPAQ  3028
dga5890.1001.0.3.0      $1$DGA5890  COMPAQ HSV110 (C)COMPAQ  3028
dga5890.1002.0.3.0      $1$DGA5890  COMPAQ HSV110 (C)COMPAQ  3028
dka0.0.0.2004.0        DKA0        COMPAQ BD03685A24  HPB7
dka100.1.0.2004.0      DKA100     COMPAQ BD01864552  3B08
dka200.2.0.2004.0      DKA200     COMPAQ BD00911934  3B00
dqa0.0.0.15.0          DQA0        HL-DT-ST CD-ROM GCR-8480  2.11
dva0.0.0.1000.0        DVA0
eia0.0.0.2005.0        EIA0        00-06-2B-03-2D-7D
pga0.0.0.3.0           PGA0        WWN 1000-0000-c92a-78e9
pka0.7.0.2004.0        PKA0        SCSI Bus ID 7
```

From the output, the interface will be EIA0 on which the IP address will be configured and can be used for cluster formation.

To obtain the interface information on Integrity server system, execute the following command on the EFI Shell:

```
Shell> fs0:
```

```
fs0:\> cd efi
```

```
fs0:\EFI> cd vms
```

```
fs0:\EFI\VMS> vms_show device
```

```
VMS: EIA0                00-30-6E-F3-EC-6E  
EFI: Acpi(HWP0002,0)/Pci(3|0)/Mac(00306EF3EC6E)
```

```
VMS: DKA100             HP 36.4GST336754LC      HPC2  
EFI: Acpi(HWP0002,100)/Pci(1|0)/Scsi(Pun1,Lun0)
```

```
VMS: DKA0               HP 36.4GMAS3367NC      HPC3      X8_3_XBJL  
EFI: fs0: Acpi(HWP0002,100)/Pci(1|0)/Scsi(Pun0,Lun0)
```

```
VMS: EWA0               00-30-6E-F3-3C-28  
EFI: Acpi(HWP0002,100)/Pci(2|0)/Mac(00306EF33C28)
```

```
fs0:\EFI\VMS>
```

From the output, the interface will be EIA0. Here fs0: is the partition of the shared system disk.

## Step 2. Executing CLUSTER\_CONFIG\_LAN.COM

Execute the following command procedure on node TULIP:

```
TULIP$ @SYS$SYSROOT:[SYSMGR]CLUSTER_CONFIG_LAN.COM;1
```

```
Cluster Configuration Procedure  
CLUSTER_CONFIG_LAN Version V2.80  
Executing on an Alpha System
```

```
DECnet Phase IV is installed on this node.  
Alpha satellites will use LANCP, not DECnet, for MOP downline  
loading.
```

```
Enter a "?" for help at any prompt.  If you are familiar with  
the execution of this procedure, you may want to mute extra notes  
and explanations by invoking it with "@CLUSTER_CONFIG_LAN BRIEF".
```

```
TULIP is an Alpha system and currently a member of a cluster  
so the following functions can be performed:
```

MAIN Menu

1. ADD an Alpha node to the cluster.

2. REMOVE a node from the cluster.
3. CHANGE a cluster member's characteristics.
4. CREATE a duplicate system disk for MRVP4.
5. MAKE a directory structure for a new root on a system disk.
6. DELETE a root from a system disk.
7. EXIT from this procedure.

Enter choice [7]: 1

This ADD function will add a new Alpha node to the cluster.

WARNING: If the node being added is a voting member, EXPECTED\_VOTES for every cluster member must be adjusted. For complete instructions check the section on configuring a cluster in the "OpenVMS Cluster Systems" manual.

CAUTION: If this cluster is running with multiple system disks and common system files will be used, please, do not proceed unless appropriate logical names are defined for cluster common files in SYLOGICALS.COM. For instructions, refer to the "OpenVMS Cluster Systems" manual.

Do you want to continue [Y]?Y

Is the node to be a clustered node with a shared SCSI/FIBRE-CHANNEL bus (Y/N)? Y

Will the node be a satellite [Y]? N

What is the node's SCS node name? JASMIN

What is the node's SCSSYSTEMID number? 14487

Will JASMIN be a boot server [Y]?Y

This procedure will now ask you for the device name of JASMIN's system root.

The default device name (DISK\$TULIPSYS:) is the logical volume name of SYS\$SYSDEVICE:.

What is the device name for JASMIN's system root  
[default DISK\$TULIPSYS:]?

What is the name of JASMIN's system root [SYS3]?SYS3

Creating directory tree SYS3 ...

**System root SYS3 created**

**ENABLE IP for cluster communications (Y/N)? Y**

**UDP port number to be used for Cluster Communication over IP[49152]?**

**Enable IP multicast for cluster communication(Y/N)[Y]?Y**

**What is the IP multicast address[224.0.0.3]?**

**What is the TTL (time to live) value for IP multicast packets [1] ?**

**Do you want to enter unicast address(es)(Y/N)[Y]?Y**

**What is the unicast address[Press [RETURN] to end the list]? 10.0.1.2**

**What is the unicast address[Press [RETURN] to end the list]? 10.0.2.2**

**What is the unicast address[Press [RETURN] to end the list]? 10.0.2.3**

**What is the unicast address[Press [RETURN] to end the list]?**

\*\*\*\*\*

Cluster Communications over IP has been enabled. Now CLUSTER\_CONFIG\_LAN will run the SYS\$MANAGER:TCPIP\$CONFIG procedure. Please select the IP interfaces to be used for Cluster Communications over IP (IPCI). This can be done selecting "Core Environment" option from the main menu followed by the "Interfaces" option. You may also use

this opportunity to configure other aspects.  
\*\*\*\*\*

Press Return to continue ...

Checking TCP/IP Services for OpenVMS configuration database files.

HP TCP/IP Services for OpenVMS Configuration Menu

Configuration options:

- 1 - Core environment
- 2 - Client components
- 3 - Server components
- 4 - Optional components
- 5 - Shutdown HP TCP/IP Services for OpenVMS
- 6 - Startup HP TCP/IP Services for OpenVMS
- 7 - Run tests
- A - Configure options 1 - 4
- [E] - Exit configuration procedure

Enter configuration option: 1

HP TCP/IP Services for OpenVMS Core Environment Configuration Menu

Configuration options:

- 1 - Domain
- 2 - Interfaces
- 3 - Routing
- 4 - BIND Resolver
- 5 - Time Zone
- A - Configure options 1 - 5
- [E] - Exit menu

Enter configuration option: 2

HP TCP/IP Services for OpenVMS Interface & Address Configuration Menu

Hostname Details: Configured=TULIP, Active=TULIP

Configuration options:

- 0 - Set The Target Node (Current Node: TULIP)
- 1 - WE0 Menu (EWA0: TwistedPair 1000mbps)
- 2 - 10.0.2.2/8 TULIP Configured,IPCI
- 3 - WE1 Menu (EWB0: TwistedPair 1000mbps)
- 4 - WE2 Menu (EWC0: TwistedPair 1000mbps)
- 5 - WE3 Menu (EWD0: TwistedPair 1000mbps)
- 6 - WE4 Menu (EWE0: TwistedPair 1000mbps)
- 7 - WE5 Menu (EWF0: TwistedPair 1000mbps)
- 8 - WE6 Menu (EWG0: Multimode 10000mbps)
- 9 - WE7 Menu (EWH0: TwistedPair 1000mbps)
- 10 - IE0 Menu (EIA0: TwistedPair 100mbps)
- 11 - IE1 Menu (EIB0: TwistedPair 100mbps)

Enter configuration option or press <ENTER> key to continue: 0 .....[1]  
Enter name of node to manage [TULIP]: JASMIN  
JASMIN is not currently a cluster member.  
\* Continue configuring JASMIN [NO]: Y.....[2]  
Enter system device for JASMIN [\$10\$DGA165:]:.....[3]  
Enter system root for JASMIN []: SYS3.....[4]

HP TCP/IP Services for OpenVMS Interface & Address Configuration Menu

Hostname Details: Configured=Not Configured

Configuration options:

- 0 - Set The Target Node (Current Node: JASMIN - \$10\$DGA165:[sys3.]
- A - Add an Interface
- [E] - Exit menu

Enter configuration option: A  
Enter controller name (e.g. EIA or EWC, etc): [ENTER when done]  
EIA.....[5]

Controller Name : EIA  
TCP/IP Interface Name : IE0

\* Is this correct [NO]: Y  
Interface Menu: IE0

HP TCP/IP Services for OpenVMS Interface IE0 Configuration Menu  
(Node: JASMIN)

Configuration options:

- 1 - Add a primary address on IE0
- 2 - Add an alias address on IE0
- 3 - Enable DHCP client to manage address on IE0
- [E] - Exit menu

Enter configuration option: 1.....[6]  
\* Is this address used by Clusters over IP (IPCI) [NO]: Y.....[7]

IPv4 Address may be entered with CIDR bits suffix.  
E.g. For a 16-bit netmask enter 10.0.1.1/16

Enter IPv4 Address []: 10.0.2.3  
Default netmask calculated from class of IP address: 255.0.0.0

IPv4 Netmask may be entered in dotted decimal notation,  
(e.g. 255.255.0.0), or as number of CIDR bits (e.g. 16)

Enter Netmask or CIDR bits [255.0.0.0]:  
Enter hostname []: JASMIN

Requested configuration:

Node : JASMIN

Interface: IE0  
IPCI : Yes  
Address : 10.0.2.3/8  
Netmask : 255.0.0.0 (CIDR bits: 8)  
Hostname : JASMIN

\* Is this correct [YES]:Y  
Added hostname JASMIN (10.0.2.3) to host database

NOTE:

The system hostname is not configured.  
It will now be set to JASMIN (10.0.2.3).  
This can be changed later via the Interface Configuration Menu.

Updated system hostname in configuration database

Added address IE0:10.0.2.3 to configuration database  
Updated Interface in IPCI configuration file:  
\$10\$DGA165:[SYS3.SYSEXEXE]TCPIP\$CLUSTER.DAT;

Updated Default Route in IPCI configuration file:  
\$10\$DGA165:[SYS3.SYSEXEXE]TCPIP\$CLUSTER.DAT;

HP TCP/IP Services for OpenVMS Interface & Address Configuration Menu

Hostname Details: Configured=JASMIN  
Configuration options:

- 0 - Set The Target Node (Current Node: JASMIN - \$10\$DGA165:[sys3.])
- 1 - IE0 Menu (EIA0:)
- 2 - 10.0.2.3/8 JASMIN Configured,IPCI
- I - Information about your configuration
- A - Add an Interface
- [E] - Exit menu

Enter configuration option:

HP TCP/IP Services for OpenVMS Core Environment Configuration Menu

Configuration options:

- 1 - Domain
- 2 - Interfaces
- 3 - Routing
- 4 - BIND Resolver
- 5 - Time Zone
- A - Configure options 1 - 5
- [E] - Exit menu

Enter configuration option:

HP TCP/IP Services for OpenVMS Configuration Menu

Configuration options:

- 1 - Core environment
- 2 - Client components

- 3 - Server components
- 4 - Optional components
- 5 - Shutdown HP TCP/IP Services for OpenVMS
- 6 - Startup HP TCP/IP Services for OpenVMS
- 7 - Run tests
- A - Configure options 1 - 4
- [E] - Exit configuration procedure

Enter configuration option:

SYS\$SYSTEM:PE\$IP\_CONFIG.DAT file generated in node JASMIN's root shown below

```
! CLUSTER_CONFIG_LAN creating for CHANGE operation on 15-JUL-2008
15:23:56.05
multicast_address=224.0.0.3
ttl=1
udp_port=49152
unicast=10.0.2.3
unicast=10.0.2.2
unicast=10.0.1.2
```

SYS\$SYSTEM:TCPIP\$CLUSTER.DAT file generated in node JASMIN's root shown below

```
interface=IE0,EIA0,10.0.2.3,255.0.0.0
default_route=16.116.40.1
```

Field	Description
[1]	In the TCP/IP configuration, select option 0 to set the target node to JASMIN, which is the new node added to the cluster.
[2]	Proceed with the configuration procedure to configure node JASMIN.
[3]	Enter the system device for JASMIN, which is \$10\$DGA165.
[4]	Enter JASMIN's root, which is SYS3.
[5]	Enter the controller information on which IP is configured for cluster traffic. This is the controller information that has been obtained from the console of the machine JASMIN as explained in the beginning of the configuration.
[6]	Select the option to add the primary address for IE0 (IP interface name of controller EIA).
[7]	Enable the use of IE0 for Cluster over IP and proceed with the rest of the configuration.

### Step 3. Completing the Configuration Procedure

Continue to run the CLUSTER\_CONFIG\_LAN.COM to complete the cluster configuration procedure, see Scenario 1 for more details.

### Step 4. Updating the PE\$IP\_CONFIG.DAT file

To ensure that the nodes join the cluster, PE\$IP\_CONFIG.DAT must be consistent through all the members of the cluster. Copy the SYS\$SYSTEM:PE\$IP\_CONFIG.DAT file that is created on node JASMIN's to the other nodes, ORCHID and TULIP.

### Step 5. Refreshing the Unicast list

On both ORCHID and TULIP, to update the new unicast list in the PE\$IP\_CONFIG.DAT file, enter the following command for PEDRIVER:

## \$MC SCACP RELOAD

You can also use SYSMAN and run the command cluster wide.

**Note:** The following rule is applicable when IP unicast address is used for node discovery. A node is allowed to join the cluster only if its IP address is present in the existing members SYS\$SYSTEM:PE\$IP\_CONFIG.DAT file.

---

### Step 6. Running AUTOGEN and Rebooting the Node

After the first boot of JASMIN, AUTOGEN.COM runs automatically. JASMIN will now be able to join the existing cluster consisting of nodes ORCHID and LOTUS.

```
JASMIN$ @SYS$UPDATE:AUTOGEN GETDATA REBOOT
```

## Scenario 4: Adding an Integrity server Satellite Node to a Cluster over IP


This section describes how to add a satellite node to an existing two-node cluster. JASMIN is an Integrity server satellite node and is added to a cluster that has two nodes, ORCHID and TULIP. TULIP is the boot server for the satellite node.

**Note:** For both Alpha and Integrity server satellite nodes, the satellite node and its boot server must exist in the same LAN segment.

---

Figure 5.0 Adding an Integrity server Satellite Node

---



## Step 1. Selecting the Interface for Satellite Booting

To select the interface to be used for satellite booting, assume that the satellite node does not have any disk running OpenVMS connected to it. Note that if you are adding Alpha systems as satellite nodes, you can get the information from ">>>" prompt by executing the following command:

```
P00>>>SHOW DEVICE
dga5245.1003.0.3.0      $1$DGA5245    COMPAQ HSV110 (C)COMPAQ  3028
dga5245.1004.0.3.0      $1$DGA5245    COMPAQ HSV110 (C)COMPAQ  3028
dga5890.1001.0.3.0      $1$DGA5890    COMPAQ HSV110 (C)COMPAQ  3028
dga5890.1002.0.3.0      $1$DGA5890    COMPAQ HSV110 (C)COMPAQ  3028
dka0.0.0.2004.0        DKA0          COMPAQ BD03685A24  HPB7
dka100.1.0.2004.0      DKA100        COMPAQ BD01864552  3B08
dka200.2.0.2004.0      DKA200        COMPAQ BD00911934  3B00
dqa0.0.0.15.0          DQA0          HL-DT-ST CD-ROM GCR-8480  2.11
dva0.0.0.1000.0        DVA0
eia0.0.0.2005.0        EIA0          00-06-2B-03-2D-7D
pga0.0.0.3.0           PGA0          WWN 1000-0000-c92a-78e9
pka0.7.0.2004.0        PKA0          SCSI Bus ID 7
pkb0.6.0.2.0           PKB0          SCSI Bus ID 6  5.57
P00>>>
```

From the output, the LAN interface will be EIA0 on which IP address will be configured and used for cluster configuration.

**Note:** The Alpha console uses the MOP protocol for network load of satellite systems. Because the MOP protocol is non-routable, the satellite boot server or servers and all satellites booting from them must reside in the same LAN. In addition, the boot server must have at least one LAN device enabled for cluster communications to permit the Alpha satellite nodes to access the system disk.

---

On Integrity server systems, the interface name either starts with 'EI' or 'EW'.

If it is the first interface, it will be EIA0 or EWA0. Note the mac address of the interface that you want to use from the Shell prompt. To obtain the interface information on Integrity servers, execute the following command on EFI Shell:

```
Shell> lanaddress
```

```
LAN Address Information
```

LAN Address	Path
Mac(00306E4A133F)	Acpi(HWP0002,0)/Pci(3 0)/Mac(00306E4A133F)
*Mac(00306E4A02F9)	Acpi(HWP0002,100)/Pci(2 0)/Mac(00306E4A02F9)

```
Shell>
```

Assuming that the interface which is active is EIA0, configure the satellite with EIA0, if it does not boot with EIA0 then try with EWA0 subsequently.

## Step 2. Executing CLUSTER\_CONFIG\_LAN.COM

Execute CLUSTER\_CONFIG\_LAN on the boot server node TULIP and select the appropriate option as described in the following example:

```
TULIP$ @SYS$SYSROOT:[SYSMGR]CLUSTER_CONFIG_LAN.COM
```

Cluster Configuration Procedure  
CLUSTER\_CONFIG\_LAN Version V2.80  
Executing on an IA64 System

DECnet-Plus is installed on this node.  
IA64 satellites will use TCP/IP BOOTP and TFTP services for downline loading

TCP/IP is installed and running on this node.

Enter a "?" for help at any prompt. If you are familiar with the execution of this procedure, you may want to mute extra notes and explanations by invoking it with "@CLUSTER\_CONFIG\_LAN BRIEF".

TULIP is an IA64 system and currently a member of a cluster so the following functions can be performed:

MAIN Menu

1. ADD an IA64 node to the cluster.
2. REMOVE a node from the cluster.
3. CHANGE a cluster member's characteristics.
4. CREATE a duplicate system disk for TULIP.
5. MAKE a directory structure for a new root on a system disk.
6. DELETE a root from a system disk.
7. EXIT from this procedure.

Enter choice [7]:

This ADD function will add a new IA64 node to the cluster.

WARNING: If the node being added is a voting member, EXPECTED\_VOTES for every cluster member must be adjusted. For complete instructions check the section on configuring a cluster in the "OpenVMS Cluster Systems" manual.

CAUTION: If this cluster is running with multiple system disks and common system files will be used, please, do not proceed unless appropriate logical names are defined for cluster common files in SYLOGICALS.COM. For instructions, refer to the "OpenVMS Cluster Systems" manual.

Do you want to continue [Y]? y

Is the node to be a clustered node with a shared SCSI/FIBRE-CHANNEL bus (Y/N)? N

Will the node be a satellite [Y]?

What is the node's SCS node name? JASMIN

What is the node's SCSSYSTEMID number? 25482

WARNING:

DECnet is not running.

No DECnet databases will be updated with information on JASMIN.

Does JASMIN need to be registered in the DECnet namespace [N]?

What is the Cluster Alias fullname?

This procedure will now ask you for the device name of JASMIN's system root.

The default device name (DISK\$TULIPSYS:) is the logical volume name of SYS\$SYSDEVICE:.

What is the device name for JASMIN's system root

[default DISK\$TULIPSYS:]?

What is the name of JASMIN's system root [SYS14]?

What is JASMIN's LAN adapter hardware address? 00-30-6E-4A-02-F9.....[1]  
What is JASMIN's TCP/IP address [10.0.2.3]? .....[2]  
What is JASMIN's TCP/IP gateway or gateways (leave blank if none)? 10.0.2.1 .....[3]  
What is JASMIN's TCP/IP network mask [255.255.255.0]? 255.255.254.0 .....[4]

NOTE: Make sure to set the VMS\_FLAGS console variable to 0,200000 on node JASMIN so it will use the memory-disk method to boot as a satellite. The command to update this variable from the console EFI shell of JASMIN is:  
set vms\_flags "0,200000"

Allow conversational bootstraps on JASMIN [N]?

The following workstation windowing options are available:

1. No workstation software
2. DECwindows Workstation Software

Enter choice [1]:

Creating directory tree SYS14 ...  
System root SYS14 created

ENABLE IP for cluster communications (Y/N)? Y .....[5]

UDP port number to be used for Cluster Communication over IP[49152]? .....[6]

Enable IP multicast for cluster communication(Y/N)[Y]? Y

What is IP the multicast address[224.0.0.3]?..... [7]

What is the TTL (time to live) value for IP multicast packets [1] ? 32]?..... [8]

Do you want to enter unicast address(es)(Y/N)[Y]? y .....[9]

What is the unicast address[Press [RETURN] to end the list]? 10.0.2.3

What is the unicast address[Press [RETURN] to end the list]? 10.0.2.2

What is the unicast address[Press [RETURN] to end the list]? 10.0.1.2

What is the unicast address[Press [RETURN] to end the list]?

```
*****
Cluster Communications over IP has been enabled. Now
CLUSTER_CONFIG_LAN will run the SYS$MANAGER:TCPIP$CONFIG
procedure. Please select the IP interfaces to be used for
Cluster Communications over IP (IPCI). This can be done
selecting "Core Environment" option from the main menu
followed by the "Interfaces" option. You may also use
this opportunity to configure other aspects.
*****
```

Press Return to continue ...

Checking TCP/IP Services for OpenVMS configuration database files.

HP TCP/IP Services for OpenVMS Configuration Menu

Configuration options:

- 1 - Core environment
- 2 - Client components
- 3 - Server components
- 4 - Optional components
- 5 - Shutdown HP TCP/IP Services for OpenVMS
- 6 - Startup HP TCP/IP Services for OpenVMS
- 7 - Run tests
- A - Configure options 1 - 4
- [E] - Exit configuration procedure

Enter configuration option: 1

HP TCP/IP Services for OpenVMS Core Environment Configuration Menu

Configuration options:

- 1 - Domain
- 2 - Interfaces
- 3 - Routing
- 4 - BIND Resolver
- 5 - Time Zone
- A - Configure options 1 - 5
- [E] - Exit menu

Enter configuration option: 2

HP TCP/IP Services for OpenVMS Interface & Address Configuration Menu

Hostname Details: Configured=[], Active=[]

Configuration options:

- 0 - Set The Target Node (Current Node: TULIP)
- 1 - IE0 Menu (EIA0: TwistedPair 100mbps)
- 2 - **10.0.2.2** \* TULIP\* Configured,Active,IPCI
- [E] - Exit menu

Enter configuration option: 0 .....[10]

Enter name of node to manage [TULIP]: JASMIN

JASMIN is not currently a cluster member.

\* Continue configuring JASMIN [NO]: YES.....[11]

Enter system device for JASMIN [DSA2:]: .....[12]

Enter system root for JASMIN []: SYS14.....[13]

HP TCP/IP Services for OpenVMS Interface & Address Configuration Menu

Hostname Details: Configured=JASMIN

Configuration options:

- 0 - Set The Target Node (Current Node: JASMIN - DSA2:[SYS14])
- A - Add an Interface
- [E] - Exit menu

Enter configuration option: A

Enter controller name (e.g. EIA or EWC, etc): [ENTER when done] EIA.....[14]

Controller Name : EIA  
TCP/IP Interface Name : IE0

\* Is this correct [NO]: y

Interface Menu:IE0

HP TCP/IP Services for OpenVMS Interface IE0 Configuration Menu (Node: JASMIN)

Configuration options:

- 1 - Add a primary address on IE0

- 2 - Add an alias address on IE0
- 3 - Enable DHCP client to manage address on IE0

[E] - Exit menu

Enter configuration option: 1.....[15]

\* Is this address used by Clusters over IP (IPCI) [NO]: y.....[16]

IPv4 Address may be entered with CIDR bits suffix.  
E.g. For a 16-bit netmask enter 10.0.1.1/16

Enter IPv4 Address []: 10.0.2.3  
Default netmask calculated from class of IP address: 255.0.0.0

IPv4 Netmask may be entered in dotted decimal notation,  
(e.g. 255.255.0.0), or as number of CIDR bits (e.g. 16)

Enter Netmask or CIDR bits [255.0.0.0]:  
Enter hostname []: JASMIN

Requested configuration:

Node : JASMIN  
Interface: IE0  
IPCI : Yes  
Address : 10.0.2.3/8  
Netmask : 255.0.0.0 (CIDR bits: 8)  
Hostname : JASMIN

\* Is this correct [YES]:  
Added hostname JASMIN (10.0.2.3) to host database

NOTE:  
The system hostname is not configured.  
It will now be set to JASMIN (10.0.2.3).  
This can be changed later via the Interface Configuration Menu.

Updated system hostname in configuration database

Added address IE1:10.0.2.3 to configuration database  
Updated Interface in IPCI configuration file:  
DSA2:[SYS14.SYSEXE]TCPIP\$CLUSTER.DAT;

Updated Default Route in IPCI configuration file:  
DSA2:[SYS14.SYSEXE]TCPIP\$CLUSTER.DAT;

### HP TCP/IP Services for OpenVMS Interface & Address Configuration Menu

Hostname Details: Configured=JASMIN

Configuration options:

- 0 - Set The Target Node (Current Node: JASMIN - DSA2:[SYS14.]
- 1 - IE1 Menu (EIB0:)
- 2 - 10.0.2.3/8 JASMIN Configured,IPCI
- I - Information about your configuration
- A - Add an Interface
- [E] - Exit menu

Enter configuration option:

HP TCP/IP Services for OpenVMS Core Environment Configuration Menu

Configuration options:

- 1 - Domain
- 2 - Interfaces
- 3 - Routing
- 4 - BIND Resolver
- 5 - Time Zone
- A - Configure options 1 - 5
- [E] - Exit menu

Enter configuration option: E

HP TCP/IP Services for OpenVMS Configuration Menu

Configuration options:

- 1 - Core environment
- 2 - Client components
- 3 - Server components
- 4 - Optional components
- 5 - Shutdown HP TCP/IP Services for OpenVMS
- 6 - Startup HP TCP/IP Services for OpenVMS
- 7 - Run tests
- A - Configure options 1 - 4
- [E] - Exit configuration procedure

Enter configuration option: E

SYSS\$SYSTEM:PE\$IP\_CONFIG.DAT file generated in node JASMIN's root shown below

```
! CLUSTER_CONFIG_LAN creating for CHANGE operation on 15-JUL-2008 15:23:56.05  
multicast_address=224.0.0.3  
ttl=32  
udp_port=49152  
unicast=10.0.2.3  
unicast=10.0.2.2  
unicast=10.0.1.2
```

SYSS\$SYSTEM:TCPIP\$CLUSTER.DAT file generated in node JASMIN's root shown below

```
interface=IE0,EIA0,10.0.2.3,255.0.0.0  
default_route=16.116.40.1
```

**Note:** Assuming that the interface which is active is EIA, configure the satellite with EIA, if it does not boot with EIA then try with EIB subsequently. If the wrong interface name is given, satellite node fails with the messages during booting.

---

Field	Description
[1]	Enter the LAN adapter's hardware address.
[2]	Enter the TCP/IP address.
[3]	Enter the TCP/IP gateway.
[4]	Enter the TCP/IP network mask address.
[5]	Enable IP for cluster communication.
[6]	The UDP port number to be used for cluster communication. The UDP port number must be same on all members of the cluster. Also, ensure that there is no other cluster in your environment using the same UDP port number and this port number must not be used by any other application.

[7]	Enter the IP multicast address for cluster, if IP multicasting is enabled. By default, the IP multicast address is selected from the administratively scoped IP multicast address range of 239.242.x.y. The last two octets x and y are generated based on the cluster group number. In the above example, the cluster group number is 1985 and can be calculated as follows: X= 1985/256 Y= 1985 - (256 *x) The System Administrator can override the default multicast address with a unique address for their environment. The multicast address is modified based on the cluster group number or it can be added to .DAT file.
[8]	TTL is the time to live for IP multicast packets. It specifies the number of hops allowed for IP multicast packets.
[9]	Enter "yes" to enter the IP Unicast address of remote nodes of the cluster, which are not reachable using IP multicast address.
[10]	In the TCP/IP configuration, select option 0 to set the target node to JASMIN, which is the satellite node, will be added to the cluster.
[11]	Proceed with configuration steps to configure node JASMIN.
[12]	Enter the system device for JASMIN, which is DSA2.
[13]	Enter JASMIN's root, which is SYS14.
[14]	Enter the controller information on which IP will be configured for cluster communication. The controller information is obtained from the console of JASMIN as explained in the beginning of the configuration.
[15]	Select an option to add a primary address for IEO (IP interface name of controller EIA).
[16]	Enable the use of IEO for Cluster over IP (IPC) and proceed with the rest of the configuration.

### Step 3. Executing the CLUSTER\_CONFIG\_LAN.COM Procedure

Continue to run the CLUSTER\_CONFIG\_LAN.COM to complete the cluster configuration procedure.

```
Adjusting protection on DSA2:[SYS14.][SYSEXE]PE$IP_CONFIG.DAT;1
Will JASMIN be a disk server [N]? Y
Enter a value for JASMIN's ALLOCLASS parameter [0]: 15
Updating BOOTP database with satellite information for JASMIN..
Size of pagefile for JASMIN [RETURN for AUTOGEN sizing]?
```

A temporary pagefile will be created until resizing by AUTOGEN. The default size below is arbitrary and may or may not be appropriate.

```
Size of temporary pagefile [10000]?
Size of swap file for JASMIN [RETURN for AUTOGEN sizing]?
```

A temporary swap file will be created until resizing by AUTOGEN. The default size below is arbitrary and may or may not be appropriate.

```
Size of temporary swap file [8000]?
```

NOTE: IA64 satellite node JASMIN requires DOSD if capturing the system state in a dumpfile is desired after a system crash.

```
Will a local disk on JASMIN be used for paging and swapping (Y/N)? N
```

If you specify a device other than DISK\$TULIPSYS: for JASMIN's page and swap files, this procedure will create PAGEFILE\_JASMIN.SYS and SWAPFILE\_JASMIN.SYS in the <SYSEXE> directory on the device you specify.

```
What is the device name for the page and swap files [DISK$TULIPSYS:]?  
%SYSGEN-I-CREATED, DSA2:<SYS14.SYSEXE>PAGEFILE.SYS;1 created  
%SYSGEN-I-CREATED, DSA2:<SYS14.SYSEXE>SWAPFILE.SYS;1 created  
The configuration procedure has completed successfully.
```

The node JASMIN is configured to join the cluster. After the first boot of JASMIN, AUTOGEN.COM will run automatically.

#### **Step 4. Updating the PE\$IP\_CONFIG.DAT File**

To ensure that the nodes join the cluster, PE\$IP\_CONFIG.DAT must be consistent through all the members of the cluster. Copy the SYS\$SYSTEM:PE\$IP\_CONFIG.DAT file that is created on node JASMIN's root to the other nodes, ORCHID and TULIP.

#### **Step 5. Refreshing the Unicast list**

On both ORCHID and TULIP, to update the new unicast list in the PE\$IP\_CONFIG.DAT file, enter the following command for PEDRIVER:

```
$MC SCACP RELOAD
```

You can also use SYSMAN and run the command cluster wide.

**Note:** The following rule is applicable when IP unicast address is used for node discovery. A node is allowed to join the cluster only if its IP address is present in the existing members of the SYS\$SYSTEM:PE\$IP\_CONFIG.DAT file.

---

#### **Step 6. Running AUTOGEN and Rebooting the Node**

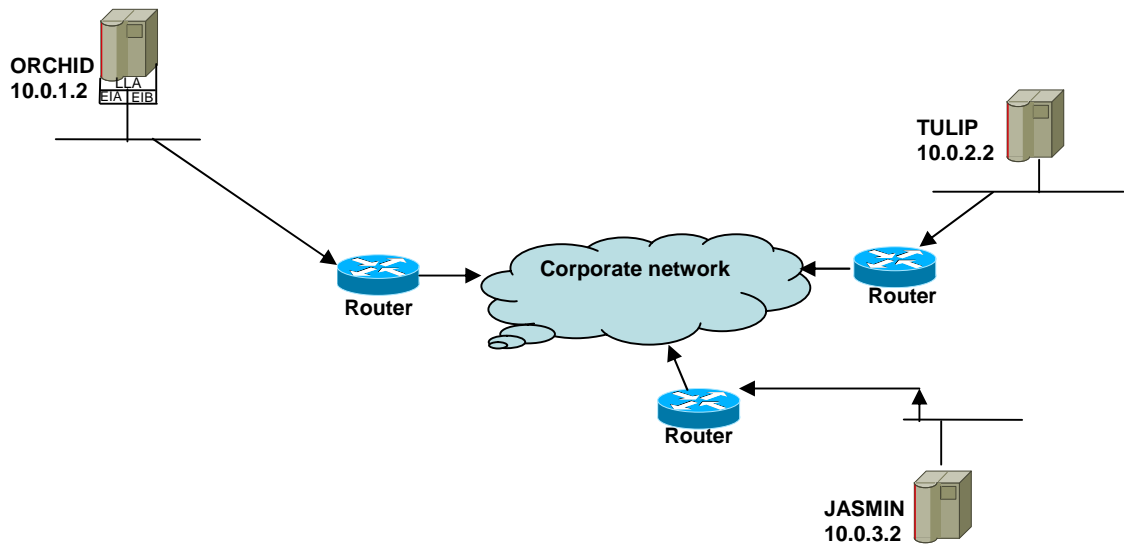
After the first boot of JASMIN, AUTOGEN.COM runs automatically. JASMIN will now be able to join the existing cluster consisting of nodes ORCHID and TULIP.

```
JASMIN$ @SYS$UPDATE:AUTOGEN GETDATA REBOOT
```

## Scenario 5: Adding an Integrity server Node to a Cluster over IP with Logical LAN Failover set

This section describes how to add a node, ORCHID to an existing two-node cluster, JASMIN and TULIP. The Logical LAN failover set is created and configured on ORCHID. ORCHID will be able to survive failure if a local LAN card fails and it will switchover to other interface configured in the logical LAN failover set.

Figure 6.0 Adding an Integrity server Node



### Step 1. Configuring the Logical LAN Failover set

Execute the following commands to create a logical LAN failover set.

```
$ MC LANCP
LANCP> DEFINE DEVICE LLB/ENABLE/FAILOVER=(EIA0, EIB0)
```

Reboot the system and during reboot, the following console messages is displayed:

```
%LLB0, Logical LAN event at 2-SEP-2008 14:52:50.06
%LLB0, Logical LAN failset device created
```

### Step 2: Executing CLUSTER\_CONFIG\_LAN

Execute CLUSTER\_CONFIG\_LAN on the node ORCHID and select the appropriate option as described in the following example:

```
ORCHID $ @SYS$MANAGER:CLUSTER_CONFIG_LAN

Cluster Configuration Procedure
CLUSTER_CONFIG_LAN Version V2.79
Executing on an IA64 System
```

DECnet-Plus is installed on this node.  
IA64 satellites will use TCP/IP BOOTP and TFTP services for downline loading.

TCP/IP is installed and running on this node.

Enter a "?" for help at any prompt. If you are familiar with the execution of this procedure, you may want to mute extra notes and explanations by invoking it with "@CLUSTER\_CONFIG\_LAN BRIEF".

This IA64 node is not currently a cluster member.

MAIN Menu

1. ADD ORCHID to existing cluster, or form a new cluster.
2. MAKE a directory structure for a new root on a system disk.
3. DELETE a root from a system disk.
4. EXIT from this procedure.

Enter choice [4]: 1

Is the node to be a clustered node with a shared SCSI/FIBRE-CHANNEL bus (Y/N)? n

What is the node's SCSI node name? ORCHID

IA64 node, using LAN/IP for cluster communications. PEDRIVER will be loaded.

No other cluster interconnects are supported for IA64 nodes.

Enter this cluster's group number: 1985

Enter this cluster's password:

Re-enter this cluster's password for verification:

ENABLE IP for cluster communications (Y/N)? Y

UDP port number to be used for Cluster Communication over IP[49152]?

Enable IP multicast for cluster communication(Y/N)[Y]? Y

What is IP the multicast address[239.242.7.193]? 239.242.7.193

What is the TTL (time to live) value for IP multicast packets [32] ?

Do you want to enter unicast address(es)(Y/N)[Y] Y

What is the unicast address[Press [RETURN] to end the list]? 10.0.1.2

What is the unicast address[Press [RETURN] to end the list]? 10.0.2.3

What is the unicast address[Press [RETURN] to end the list]? 10.0.2.2

What is the unicast address[Press [RETURN] to end the list]?

```
*****
Cluster Communications over IP has been enabled. Now
CLUSTER_CONFIG_LAN will run the SYS$MANAGER:TCPIP$CONFIG
procedure. Please select the IP interfaces to be used for
Cluster Communications over IP (IPCI). This can be done
selecting "Core Environment" option from the main menu
followed by the "Interfaces" option. You may also use
this opportunity to configure other aspects.
*****
```

Press Return to continue ...

#### TCP/IP Network Configuration Procedure

This procedure helps you define the parameters required to run HP TCP/IP Services for OpenVMS on this system.

%TCPIP-I-IPCI, TCP/IP Configuration is limited to IPCI.  
-TCPIP-I-IPCI, Rerun TCPIP\$CONFIG after joining the cluster.

HP TCP/IP Services for OpenVMS Interface & Address Configuration Menu

Hostname Details: Configured=Not Configured, Active=nodeg

Configuration options:

- 0 - Set The Target Node (Current Node: ORCHID)
- 1 - LE0 Menu (LLA0: TwistedPair 100mbps)
- 2 - IE1 Menu (EIB0: TwistedPair 100mbps)
- [E] - Exit menu

Enter configuration option: 1

\* IPCI Address Configuration \*

Only IPCI addresses can be configured in the current environment.  
After configuring your IPCI address(es) it will be necessary to  
run TCPIP\$CONFIG once your node has joined the cluster.

IPv4 Address may be entered with CIDR bits suffix.  
E.g. For a 16-bit netmask enter 10.0.1.1/16

Enter IPv4 Address []:10.0.1.2

Default netmask calculated from class of IP address: 255.0.0.0

IPv4 Netmask may be entered in dotted decimal notation,  
(e.g. 255.255.0.0), or as number of CIDR bits (e.g. 16)

Enter Netmask or CIDR bits [255.0.0.0]: 255.255.255.0

Requested configuration:

Node : ORCHID  
Interface: IE0  
IPCI : Yes  
Address : 10.0.1.2/24  
Netmask : 255.255.254.0 (CIDR bits: 23)

\* Is this correct [YES]:

Updated Interface in IPCI configuration file:

SY\$SYSROOT:[SYSEXE]TCPIP\$CLUSTER.

DAT;

HP TCP/IP Services for OpenVMS Interface & Address Configuration Menu

Hostname Details: Configured=Not Configured, Active=nodeg

Configuration options:

- 0 - Set The Target Node (Current Node: ORCHID)
- 1 - LE0 Menu (LLA0: TwistedPair 100mbps)
- 2 - 10.0.1.2/24 ORCHID IPCI

```
3 - IE1 Menu (EIB0: TwistedPair 100mbps)
[E] - Exit menu
```

```
Enter configuration option: E
Enter your Default Gateway address []: 10.0.1.1
* The default gateway will be: 10.0.1.1 Correct [NO]: YES
Updated Default Route in IPCI configuration file:
SYS$SYSROOT:[SYSEXE]TCPIP$CLUSTER.DAT;
TCPIP-I-IPCIDONE, Finished configuring IPCI address information.
```

SYS\$SYSTEM:PE\$IP\_CONFIG.DAT file generated in node ORCHID's root shown below

```
! CLUSTER_CONFIG_LAN creating for CHANGE operation on 15-JUL-2008
15:23:56.05
multicast_address=239.242.7.193
ttl=32
udp_port=49152
unicast=10.0.2.3
unicast=10.0.2.2
unicast=10.0.1.2
```

SYS\$SYSTEM:TCPIP\$CLUSTER.DAT file generated in node ORCHID's root shown below

```
interface=LE1,LLB0,10.0.1.2,255.0.0.0
default_route=10.0.1.1
```

### **Step 3. Completing the Configuration Procedure**

Continue to run the CLUSTER\_CONFIG\_LAN.COM to complete the cluster configuration procedure. For more information, see Scenario 1.

### **Step 4. Updating the PE\$IP\_CONFIG.DAT file**

To ensure that the nodes join the cluster, PE\$IP\_CONFIG.DAT must be consistent through all the members of the cluster. Copy the SYS\$SYSTEM:PE\$IP\_CONFIG.DAT file that is created on node JASMIN to the other nodes ORCHID and TULIP.

### **Step 5. Refreshing the Unicast list**

On both ORCHID and TULIP, to update the new unicast list in the PE\$IP\_CONFIG.DAT file, enter the following command for PEDRIVER:

```
$MC SCACP RELOAD
```

You can also use SYSMAN and run the command cluster wide.

### **Step 6. Running AUTOGEN and Rebooting the Node**

After the first boot of ORCHID, AUTOGEN.COM will run automatically. ORCHID will now be able to join the existing cluster consisting of nodes JASMIN and LOTUS.

```
ORCHID$ @SYS$UPDATE:AUTOGEN GETDATA REBOOT
```

# Appendix A - Cluster over IP Commands

This section describes the commands used to configure and monitor cluster systems interconnected with IP.

## SHOW CHANNEL /ALL

Displays the PEdriver channel information for all the nodes. Each channel is a single NISCA communication path between a LAN device on a local system and a LAN device on a remote system. Use the SHOW CHANNEL command to display the node names that are local and remote device names.

Example:

```
SCACP> SHOW CHANNEL/ALL
```

```
MADDY$ MC SCACP SHOW CHANNEL/ALL
```

```
MADDY PEA0 Channel Summary 15-FEB-2009 12:08:24.16:
```

Remote Node Time	Device Loc Rmt	Channel State	Total Errors	ECS State	Priority Cur Mgt	Buffer Hops Size	Delay (uSec)	Load Class	Total Pkts(S+R)	----- Most Recent ----- CH Opened Time	CH Closed
ORCHID	WE0 IE0	Open	57	Y(T,P,F)	0 0	2 1394	4198.4	100	1831428	14-FEB 11:35:11.54	(No time)
TULIP	WE0 IE0	Open	1	N(T,I,F)	0 0	2 1394	250.0	100	196042	14-FEB 11:35:12.05	(No time)
TULIP	EIA EIA	Open	1	Y(T,P,F)	0 0	2 1426	496.2	100	634047	14-FEB 11:34:46.92	(No time)
TULIP	EIA EIB	Open	1	Y(T,P,F)	0 0	2 1426	546.0	100	444646	14-FEB 11:34:45.04	(No time)
TULIP	EWA EIA	Open	0	Y(T,P,F)	0 0	2 1426	521.6	1000	675722	14-FEB 11:34:38.06	(No time)
TULIP	EWA EIB	Open	1	Y(T,P,F)	0 0	2 1426	514.9	1000	503848	14-FEB 11:34:38.06	(No time)
JASMIN	WE0 IE1	Open	1	N(T,I,F)	0 0	2 1394	250.0	100	154648	14-FEB 11:35:12.10	(No time)
JASMIN	EIA EIA	Open	1	Y(T,P,F)	0 0	2 1426	455.7	100	471521	14-FEB 11:34:46.04	(No time)
JASMIN	EIA EIB	Open	1	Y(T,P,F)	0 0	2 1426	450.3	100	434176	14-FEB 11:34:45.04	(No time)
JASMIN	EWA EIB	Open	1	Y(T,P,F)	0 0	2 1426	457.4	100	471021	14-FEB 11:34:39.25	(No time)
JASMIN	EWA EIA	Open	1	Y(T,P,F)	0 0	2 1426	423.7	100	547601	14-FEB 11:34:38.84	(No time)
MADDY	EWA EIA	Open	1	Y(T,P,F)	0 0	2 1426	250.0	100	156529	14-FEB 11:34:45.04	(No time)
MADDY	EIA EWA	Open	1	Y(T,P,F)	0 0	2 1426	250.0	100	156218	14-FEB 11:34:45.04	(No time)
MADDY	LCL LCL	Open	3	Y(T,P,F)	0 0	2 1426	246.9	0	155833	14-FEB 11:34:38.62	(No time)

```
MADDY PEA0 Channel Miscellaneous 15-FEB-2009 12:08:24.16:
```

Remote Node	LAN/IP Local	Device Remote	Local LAN/IP Address	Remote LAN/IP Address	Cur	Neg	Mgt	Loc	Rmt	Change	Local	Remote
ORCHID	WE0	IE0	15.146.235.10	16.138.182.6	1394	1394	8120	1394	1394	1	1	2
TULIP	WE0	IE0	15.146.235.10	15.146.235.222	1394	1394	8120	1394	1394	1	1	2
TULIP	EIA 82559	EIA	00-30-6E-4A-13-3F	00-17-08-7C-B9-16	1426	1426	8120	1426	1426	1	1	2
TULIP	EIA 82559	EIB	00-30-6E-4A-13-3F	00-17-08-7C-B9-17	1426	1426	8120	1426	1426	1	1	2
TULIP	EWA	EIA	00-30-6E-4A-02-F9	00-17-08-7C-B9-16	1426	1426	8120	1426	1426	1	1	2
TULIP	EWA	EIB	00-30-6E-4A-02-F9	00-17-08-7C-B9-17	1426	1426	8120	1426	1426	1	1	2
JASMIN	WE0	IE1	15.146.235.10	15.146.235.19	1394	1394	8120	1394	1394	1	1	2
JASMIN	EIA 82559	EIA 82558	00-30-6E-4A-13-3F	00-08-02-3E-92-10	1426	1426	8120	1426	1426	1	1	2
JASMIN	EIA 82559	EIB 82558	00-30-6E-4A-13-3F	00-08-02-3E-92-11	1426	1426	8120	1426	1426	1	1	2
JASMIN	EWA	EIB 82558	00-30-6E-4A-02-F9	00-08-02-3E-92-11	1426	1426	8120	1426	1426	1	1	2
JASMIN	EWA	EIA 82558	00-30-6E-4A-02-F9	00-08-02-3E-92-10	1426	1426	8120	1426	1426	1	1	2
MADDY	EWA	EIA 82559	00-30-6E-4A-02-F9	00-30-6E-4A-13-3F	1426	1426	8120	1426	1426	1	1	1
MADDY	EIA 82559	EWA	00-30-6E-4A-13-3F	00-30-6E-4A-02-F9	1426	1426	8120	1426	1426	1	1	1
MADDY	LCL	LCL	00-00-00-00-00-00	00-00-00-00-00-00	1426	1426	8120	1426	1426	1	1	1

```
MADDY PEA0 Channel Equivalent Channel Set (ECS) 15-FEB-2009 12:08:24.16:
```

Remote Node	Device Loc Rmt	ECS State	ECS Losses	ECS Transitions	Average RTTime	Remote Ring Size	Remote Rcv Cache	Time in Open State	ECS Member Time
ORCHID	WE0 IE0	Y(T,P,F)	0	19	4198.4	8	32	00:33:12.62	00:33:12.62
TULIP	WE0 IE0	N(T,I,F)	0	0	250.0	8	64	00:33:12.10	(No time)
TULIP	EIA EIA	Y(T,P,F)	0	2191	496.2	128	64	00:33:37.23	18:34:24.66
TULIP	EIA EIB	Y(T,P,F)	0	2107	546.0	128	64	00:33:39.11	18:37:21.38
TULIP	EWA EIA	Y(T,P,F)	0	2153	521.6	128	64	00:33:46.09	18:59:29.32
TULIP	EWA EIB	Y(T,P,F)	0	2243	514.9	128	64	00:33:46.09	18:32:49.40
JASMIN	WE0 IE1	N(T,I,F)	0	0	250.0	8	32	00:33:12.05	(No time)
JASMIN	EIA EIA	Y(T,P,F)	0	1541	455.7	32	32	00:33:38.11	21:47:05.88
JASMIN	EIA EIB	Y(T,P,F)	0	1511	450.3	32	32	00:33:39.11	21:53:49.76
JASMIN	EWA EIB	Y(T,P,F)	0	1417	457.4	32	32	00:33:44.91	21:54:41.59

JASMIN	EWA	EIA	Y(T,P,F)	0	1359	423.7	32	32	00:33:45.32	21:56:30.43
MADDY	EWA	EIA	Y(T,P,F)	0	3	250.0	32	32	00:33:39.11	00:32:39.68
MADDY	EIA	EWA	Y(T,P,F)	0	3	250.0	128	32	00:33:39.11	00:32:43.68
MADDY	LCL	LCL	Y(T,P,F)	0	1	246.9	8	32	00:33:45.53	00:33:45.53

MADDY PEA0 Channel Counters and Errors 15-FEB-2009 12:08:24.16:

Remote Node	Device Loc Rmt	-- Transmit --		-- Receive --		Xmt:Rexmit	Rexmit Errors	TransmitFail Penalties	Receive Errors	Other Errors
		Messages	Bytes	Messages	Bytes					
ORCHID	WE0 IE0	805782	142672600	1025646	430448532	14389	56	0	1	0
TULIP	WE0 IE0	58927	6129756	137115	14261308	58927	1	0	0	0
TULIP	EIA EIA	236620	45342464	397427	184255260	236620	1	0	0	0
TULIP	EIA EIB	229440	43549741	215206	71208495	229440	1	0	0	0
TULIP	EWA EIA	244550	46781318	431172	181262983	Infinite	0	0	0	0
TULIP	EWA EIB	225133	42441919	278715	121749892	225133	1	0	0	0
JASMIN	WE0 IE1	58198	6053940	96450	10032148	58198	1	0	0	0
JASMIN	EIA EIA	212355	39590192	259166	90998564	212355	1	0	0	0
JASMIN	EIA EIB	196205	36172035	237971	80468129	196205	1	0	0	0
JASMIN	EWA EIB	212469	39677191	258552	90156844	212469	1	0	0	0
JASMIN	EWA EIA	213187	40135661	334414	128234144	213187	1	0	0	0
MADDY	EWA EIA	58944	6131556	97585	10150220	58944	1	0	0	0
MADDY	EIA EWA	58945	6131660	97273	10117772	58945	1	0	0	0
MADDY	LCL LCL	58956	6132702	96877	10076486	Infinite	0	0	0	3

MADDY PEA0 Channel Errors 15-FEB-2009 12:08:24.16:

Remote Node	Device Loc Rmt	Receive						-- Timeouts --		No MSCP Server	Disk Not Served	Path Restart
		Bad	Auth	Bad ECO	Bad MCA	Short	Incompat	Old Chan	Handshake Listen			
ORCHID	WE0 IE0	1	0	0	0	0	0	0	0	0	0	0
TULIP	WE0 IE0	0	0	0	0	0	0	0	0	0	0	0
TULIP	EIA EIA	0	0	0	0	0	0	0	0	0	0	0
TULIP	EIA EIB	0	0	0	0	0	0	0	0	0	0	0
TULIP	EWA EIA	0	0	0	0	0	0	0	0	0	0	0
TULIP	EWA EIB	0	0	0	0	0	0	0	0	0	0	0
JASMIN	WE0 IE1	0	0	0	0	0	0	0	0	0	0	0
JASMIN	EIA EIA	0	0	0	0	0	0	0	0	0	0	0
JASMIN	EIA EIB	0	0	0	0	0	0	0	0	0	0	0
JASMIN	EWA EIB	0	0	0	0	0	0	0	0	0	0	0
JASMIN	EWA EIA	0	0	0	0	0	0	0	0	0	0	0
MADDY	EWA EIA	0	0	0	0	0	0	0	0	0	0	0
MADDY	EIA EWA	0	0	0	0	0	0	0	0	0	0	0
MADDY	LCL LCL	0	0	0	0	0	0	0	0	0	3	0

Data	Description
Seq Retransmit	Number of times a sequenced VC packet sent on this channel was retransmitted, and the channel was penalized for the lost packet. Note that the sequential retransmit is not necessarily a reflection of lost packet. It is possible that there can be a PE which could have triggered a retransmitted and results in a duplicate packet to be sent. This is reflected in the number of duplicate packets received in the remote node. The XMIT:REXMT ratio is also a measure of for how many transmitted packet, a packet was retransmitted. A very low value (less than 1000) reflects possible network congestion.
LAN Transmit Failures	Number of times the local LAN device reported a failure to transmit a packet, and channel was penalized for the lost packet.
Restart Channel	Close or restart because channel control packet received indicating that the other end closed the channel and is restarting the channel handshake.
Channel Init Timeouts	Channel initialization handshake timeout.
Listen Timeouts	No packets of any kind, including HELLOs, were received in LISTEN_TIMEOUT seconds.
Bad Authorization Msg	Received a channel control (CC) packet with a bad authorization field.

Bad ECO CC Msg	Received a CC packet with an incompatible NISCA protocol ECO rev. field value.
Bad Multicast Msg	Received a bad multicast CC packet.
CC Short Packet	Received a CC packet that was short.
CC Incompatible	Received a CC packet that was incompatible with existing channels for this virtual circuit.
Rcv Old Channel	Received a packet from an old instance of a channel.
No MSCP Server	No MSCP server available to respond to a received channel control solicit service packet asking this node to boot serve another node.
Disk Not Served	Disk is not served by this system.
Buffer Size Change	Change in buffer size.

### ECS State Channel ECS membership information:

- Y is a member
- N is a non-member
- Losses is one of the following:
  - T (tight) - Packet loss history is acceptable.
  - L (lossy) - Recent history of packet losses makes channel unusable.
- Capacity is one of the following:
  - P (peer) - Priority and Buffer size both match the highest corresponding values of the set of tight channels, entitling the channel to be an ECS member.
  - I (inferior) - Priority or buffer size does not match the corresponding values of the set of tight channels.
  - S (superior) - Priority or buffer size is better than those of the current corresponding values of the set ECS member channels. This is a short-lived, transient state because it exists only while the ECS membership criteria are being re-evaluated.
  - U (unevaluated) - Priority or buffer size, or both, have not been evaluated against the ECS criteria, usually because the channel is lossy.
- Speed is one of the following:
  - F (fast) - Channel delay is among the best for tight and peer channels.

### SHOW CHANNEL /IP

This command displays the IP channel data summary.

Example:

MADDY\$SHOW CHANNEL/IP

MADDY PEA0 Channel Summary 15-FEB-2009 11:44:56.92:

Remote Node	IP Loc	Dev Rmt	Channel State	Total Errors	ECS State	Priority Cur	Priority Mgt	Hops	Buffer Size	Delay (uSec)	Load Class	Total Pkts(S+R)	CH Opened	Most Recent Time	CH Closed	Time
ORCHID	WE0	IE0	Open	38	Y(T,P,F)	0	0	2	1394	4975.0	100	521626	14-FEB	11:35:11.54	(No time)	
TULIP	WE0	IE0	Open	1	N(T,I,F)	0	0	2	1394	250.0	100	192922	14-FEB	11:35:12.05	(No time)	
JASMIN	WE0	IE1	Open	1	N(T,I,F)	0	0	2	1394	250.0	100	152181	14-FEB	11:35:12.10	(No time)	

## SHOW CHANNEL /LANCHANNEL

SHOW CHANNEL/LANCHANNEL displays the LAN channel data summary.

Example:

```
$MC SCACP SHOW CHANNEL/LANCHANNEL
```

```
MADDY PEA0 Channel Summary 15-FEB-2009 11:37:07.13:
```

Remote Node	LAN Loc	Dev Rmt	Channel State	Total Errors	ECS State	Priority Cur	Mgt Hops	Buffer Size	Delay (uSec)	Load Class	Total Pkts(S+R)	CH Opened Time	---- Most Recent CH Closed Time
TULIP	EIA	EIA	Open	1	Y(T,P,F)	0	0	2 1426	2631.3	100	201210	14-FEB 11:34:46.92	(No time)
TULIP	EIA	EIB	Open	1	Y(T,P,F)	0	0	2 1426	2880.1	100	178304	14-FEB 11:34:45.04	(No time)
TULIP	EWA	EIA	Open	0	Y(T,P,F)	0	0	2 1426	2818.4	1000	189590	14-FEB 11:34:38.06	(No time)
TULIP	EWA	EIB	Open	1	N(T,P,S)	0	0	2 1426	3146.7	1000	209654	14-FEB 11:34:38.06	(No time)
JASMIN	EIA	EIA	Open	1	Y(T,P,F)	0	0	2 1426	4454.9	100	157760	14-FEB 11:34:46.04	(No time)
JASMIN	EIA	EIB	Open	1	Y(T,P,F)	0	0	2 1426	4332.5	100	157196	14-FEB 11:34:45.04	(No time)
JASMIN	EWA	EIB	Open	1	Y(T,P,F)	0	0	2 1426	4682.1	100	158289	14-FEB 11:34:39.25	(No time)
JASMIN	EWA	EIA	Open	1	Y(T,P,F)	0	0	2 1426	4967.5	100	159365	14-FEB 11:34:38.84	(No time)
MADDY	EWA	EIA	Open	1	Y(T,P,F)	0	0	2 1426	250.0	100	153210	14-FEB 11:34:45.04	(No time)
MADDY	EIA	EWA	Open	1	Y(T,P,F)	0	0	2 1426	250.0	100	152894	14-FEB 11:34:45.04	(No time)
MADDY	LCL	LCL	Open	3	Y(T,P,F)	0	0	2 1426	246.9	0	152524	14-FEB 11:34:38.62	(No time)

### ECS State Channel ECS membership information:

- Y is a member
- N is a non-member
- Losses is one of the following:
  - T (tight) - Packet loss history is acceptable.
  - L (lossy) - Recent history of packet losses makes channel unusable.
- Capacity is one of the following:
  - P (peer) - Priority and Buffer size both match the highest corresponding values of the set of tight channels, entitling the channel to be an ECS member.
  - I (inferior) - Priority or buffer size does not match the corresponding values of the set of tight channels.
  - S (superior) - Priority or buffer size is better than those of the current corresponding values of the set ECS member channels. This is a short-lived, transient state because it exists only while the ECS membership criteria are being re-evaluated.
  - U (unevaluated) - Priority or buffer size, or both, have not been evaluated against the ECS criteria, usually because the channel is lossy.
- Speed is one of the following:
  - F (fast) - Channel delay is among the best for tight and peer channels.

## SHOW CHANNEL /ECS

SHOW CHANNEL /ECS displays the ECS state of the channel

Example:

```
$MC SCACP SHOW CHANNEL/ECS
```

```
MADDY PEA0 Channel Summary 15-FEB-2009 11:44:02.02:
```

Remote Node	Device Loc	Rmt	Channel State	Total Errors	ECS State	Priority Cur	Mgt Hops	Buffer Size	Delay (uSec)	Load Class	Total Pkts(S+R)	CH Opened Time	---- Most Recent CH Closed Time
ORCHID	WE0	IE0	Open	38	Y(T,P,F)	0	0	2 1394	3945.8	100	470836	14-FEB 11:35:11.54	(No time)
TULIP	EIA	EIA	Open	1	Y(T,P,F)	0	0	2 1426	503.4	100	269393	14-FEB 11:34:46.92	(No time)
TULIP	EIA	EIB	Open	1	Y(T,P,F)	0	0	2 1426	459.5	100	213800	14-FEB 11:34:45.04	(No time)
TULIP	EWA	EIA	Open	0	Y(T,P,F)	0	0	2 1426	574.3	1000	284950	14-FEB 11:34:38.06	(No time)
TULIP	EWA	EIB	Open	1	Y(T,P,F)	0	0	2 1426	477.7	1000	245966	14-FEB 11:34:38.06	(No time)
JASMIN	EIA	EIA	Open	1	Y(T,P,F)	0	0	2 1426	446.9	100	227512	14-FEB 11:34:46.04	(No time)
JASMIN	EIA	EIB	Open	1	Y(T,P,F)	0	0	2 1426	428.6	100	181574	14-FEB 11:34:45.04	(No time)

JASMIN	EWA	EIB	Open	1	Y(T,P,F)	0	0	2	1426	490.5	100	219178	14-FEB	11:34:39.25	(No time)
JASMIN	EWA	EIA	Open	1	Y(T,P,F)	0	0	2	1426	458.4	100	209228	14-FEB	11:34:38.84	(No time)
MADDY	EWA	EIA	Open	1	Y(T,P,F)	0	0	2	1426	250.0	100	153948	14-FEB	11:34:45.04	(No time)
MADDY	EIA	EWA	Open	1	Y(T,P,F)	0	0	2	1426	250.0	100	153626	14-FEB	11:34:45.04	(No time)
MADDY	LCL	LCL	Open	3	Y(T,P,F)	0	0	2	1426	246.9	0	153257	14-FEB	11:34:38.62	(No time)

## ECS State Channel ECS Membership Information

OpenVMS uses multiple interfaces to communicate with any other node in order to do load balancing of communication. However, at a given time not all interfaces that link the remote node are used to transmit datagrams.

OpenVMS maintains a set of equivalent channels Equivalent Channel Set (ECS) within a VC. These channels have approximately equivalent transmission quality at a given time. Only the channels within the ECS are used to transmit datagrams to the given node. "A" is the generic format and may be "Y" (Yes) or "N" (No) stating whether the channel is in ECS or not. The remaining characters specify the quality of the channel as they are derived from the channel performance data. The characters are:

- A: T or L for Tight or Lossy
- B: P, S, I, U for Peer, Superior, Inferior or Ungraded
- C: F or S for Fast or Slow

For more details about ECS, see the section NISCA Transport Protocol Channel Selection and Congestion Control in the *HP OpenVMS Cluster Systems* manual.

**Note:** From OpenVMS Version 8.3 onwards, Topology change column from SHOW CHANNEL/FULL or /5 has been removed. Because you must not consider this as an 'error' as it is the count of failovers from one interconnect to the other interconnect. Whenever a failover occurs to another interconnect the buffer size changes. Hence this topology change is counted under "Buffer SizeDecr" column in SHOW VC/FULL output.

## SHOW IP\_INTERFACE

Displays the PEdriver device IP interface data. Each device is an IP interface on the system, which can be used for NISCA communication.

### Format

SHOW IP\_INTERFACE <ipinterfacename>

\$MC SCACP SHOW IP\_INTERFACE/ALL

MADDY PEA0 Device Summary 15-FEB-2009 11:43:35.00:

Device	Type	Errors + Events	Status	Mgt Priority	Buffer Size	MgtMax BufSiz	Line Speed	Total Pkts(S+R)	Current IP Address	
WE0		0	Run Online	XChain_Disabled	0	1394	0	N/A	945220	15.146.235.10

MADDY PEA0 Device Counters and Errors 15-FEB-2009 11:43:35.00:

Device	Messages	Transmit Bytes	Errors	Receive Messages	Receive Bytes	Errors	Port Usable	Port Down	Addr Change	Restart Failed	*Last Event Time	
WE0	541783	72155739	0	403437	93906139	0	0	1	0	0	0	15-FEB-11:43:34.99

MADDY PEA0 Device Errors 15-FEB-2009 11:43:35.00:

Device	Generic Errors	Transmit Hello Errors	Last Transmit Error	With Bad SYSTEMID	Mcast Message	TR Message	Short CC Msg	Short DX Msg	On Wrong Port	Discard PortDisa	Allocation Failures On CH	On VC
WE0	0	0	(No time)	0	0	0	0	0	0	0	0	0

Data	Description
Bad SCSSYSTEM ID	Received a packet with the wrong SCSSYSTEM ID in it.
MC Msgs Directed to TR Layer	Number of multicast packets directed to the NISCA Transport layer.
Short CC Messages Received	Number of packets received that were short to contain a NISCA channel control header.
Short DX Messages Received	Number of packets received that were short to contain a NISCA DX header for a LAN device.
CH Allocation Failures	Number of times the system failed to allocate memory for use as a channel structure in response to a packet received by this LAN device.
VC Allocation Failures	Number of times the system failed to allocate memory for use as a VC structure in response to a packet received by this LAN or IP device.
Wrong Port	Number of packets addressed to the wrong NISCA address (Invalid cluster number).
Port Disabled	Number of packets discarded because the LAN or IP device was disabled.
H/W Transmit Errors	Number of local hardware transmit errors.
Hello Transmit Errors	Number of transmit errors during HELLOs.
Last Transmit Error	Reason Reason for last transmit error.
Time of Last Transmit Error	Time of last transmit error: date and time.

### SHOW IP\_INTERFACE/COUNTERS

This command displays IP interface summary data. By default, /SUMMARY is considered if /ALL, /COUNTERS, and /SDA qualifiers are not specified.

#### Example

```
SCACP> SHOW IP_INTERFACE/COUNTERS
```

The following command displays IP interface counters.

```
SCACP> SHOW IP_INTERFACE/COUNTERS
```

```
BEEBLE PEA0 Device Counters and Errors 20-AUG-2008 14:56:29.27:
```

Device	----- Transmit -----			----- Receive -----			Port Usable	Port Down	Addr Change	Restart Failed	*Last Event Time
	Messages	Bytes	Errors	Messages	Bytes	Errors					
IE1	1229900	154289219	0	3061620	803269199	0	1	0	0	0	20-AUG 14:56:28.42
IE0	16020765	3591210844	0	3389450	879201501	0	1	0	0	0	20-AUG 14:56:27.76

```
BEEBLE PEA0 Device Errors 20-AUG-2008 14:56:29.27:
```

Device	-- Transmit --			----- Received -----						Allocation		
	Generic Errors	Hello Errors	Last Transmit Error	With Bad SYSTEMID	Mcast Message	TR	Short CC Msg	Short DX Msg	On Wrong Port	Discard PortDisa	Failures CH	On VC
IE1	0	0	(No time)	0	0	0	0	0	0	0	0	0
IE0	0	0	(No time)	0	0	0	0	0	0	0	0	0

```
SCACP> SHOW IP_INTERFACE/COUNTERS/INTERVAL
```

```
SCACP> SPAWN WAIT 0:0:10
```

```
SCACP> SHOW IP_INTERFACE/COUNTERS/INTERVAL
```

The first command in this example displays IP interface counters since the last SHOW command was executed. The SPAWN command requests the DCLWAIT command to insert a 10-second delay. The second SHOW command displays counters after a delay of 10 seconds.

## Using SDA to Monitor Cluster Communications

If your system shows symptoms of intermittent failures during run time, you need to determine whether there is a network problem or whether the symptoms are caused by some other activity in the system. Generally, you can diagnose problems in the NISCA protocol or the network using the OpenVMS System Dump Analyzer utility (SDA). SDA is an effective tool for isolating problems on specific nodes running in the OpenVMS Cluster system.

### **SDA Command SHOW PORT**

The SDA command SHOW PORT provides relevant information that is useful for troubleshooting PEDRIVER and LAN adapters in particular. Begin by entering the SHOW PORT command, which causes SDA to define cluster symbols.

The following example illustrates how the SHOW PORT command provides a summary of OpenVMS Cluster data structures.

#### **Example - SDA Command SHOW PORT Display**

```
$ ANALYZE/SYSTEM
SDA> SHOW PORT
VAXcluster data structures
-----
--- PDT Summary Page ---
PDT Address Type Device Driver Name
-----
80C3DBA0 pa PAA0 PADRIVER
80C6F7A0 pe PEA0 PEDRIVER
```

### **Monitoring Virtual Circuits**

To examine information about the virtual circuit that carries messages between the local node (where you are running SDA) and another remote node, enter the SDA command SHOW PORT/VC=VC\_remote-node-name. The following example shows how to examine information about the virtual channel running between a local node and the remote node NODE11.

#### **Example - SDA Command SHOW PORT/VC Display**

```
SDA> SHOW PORT/VC=VC_NODE11
```

### **Monitoring PEDRIVER for LAN devices**

The SDA command PE LAN\_DEVICE is useful for displaying PEDRIVER LAN device data. Each LAN device is a local LAN device on the system being used for NISCACP communications.

```
SDA> PE LAN_DEVICE
```

In the following example PE LAN\_DEVICE displays the LAN device summary of I64MOZ

## Example - SDA Command PE LAN\_DEVICE Display

```
SDA> PE LAN_DEVICE

PE$SDA Extension on I64MOZ (HP rx4640 (1.50GHz/6.0MB)) at 13-OCT-2008 14:45:45.31
-----
--

I64MOZ Device Summary 13-OCT-2008 14:45:45.31:

      Device Line Buffer MgtBuf Load Mgt Current Total
Errors & Device Type Speed Size SizeCap Class Priority LAN Address Bytes
Events Status
-----
LCL 0 Run Online Local Restart 0 0 0 00-00-00-00-00-00 120148140
EIA 2 Run Online Restart 100 1426 0 1000 0 00-30-6E-5D-97-AE 19621754
EIB 884500 Run Online Restart 1000 1426 0 1000 0 00-30-6E-5D-97-AF 0
```

### Monitoring PEDRIVER Buses for LAN devices

The SDA command SHOW PORT/BUS=BUS\_LAN-device command is useful for displaying the PEDRIVER representation of a LAN adapter. To PEDRIVER, a bus is the logical representation of the LAN adapter. To list the names and addresses of buses, enter the SDA command SHOW PORT/ADDR=PE\_PDT and then press the Return key twice. The following example shows a display for the LAN adapter named EXA.

### Example - SDA Command SHOW PORT/BUS Display

```
SDA> SHOW PORT/BUS=BUS_EXA
```

### Monitoring LAN Adapters

Use the SDA command SHOW LAN/COUNT to display the information about the LAN adapters as maintained by the LAN device driver. The command shows counters for all protocols, and not just PEDRIVER [SCA] related counters. The following example shows a sample display from the SHOW LAN/COUNTERS command.

### Example - SDA Command SHOW LAN/COUNTERS Display

```
$ ANALYZE/SYSTEM
SDA> SHOW LAN/COUNTERS
```

Use the SDA command SHOW LAN/ALL to display all the information about LAN adapters. The following example shows a sample display from the SHOW LAN/ALL command.

### Example SDA Command SHOW LAN/ALL Display

```
$MC SCACP SHOW LAN/ALL
```

```
MADDY PEA0 Device Summary 15-FEB-2009 11:43:12.89:
```

Device	Device Type	Errors + Events	Status	Mgt Priority	Buffer Size	MgtMax BufSiz	Line Speed	Total Pkts(S+R)	Current LAN Address
LCL		0	Run Online Local Restart	0	1426	0	N/A	190438	00-00-00-00-00-00
EWA		0	Run Online Restart	0	1426	0	100	1114327	00-30-6E-4A-02-F9
EIA	82559	0	Run Online Restart	0	1426	0	100	1041011	00-30-6E-4A-13-3F

```
MADDY PEA0 Device Counters and Errors 15-FEB-2009 11:43:12.89:
```

Device	Messages	Transmit Bytes	Errors	Receive Messages	Bytes	Errors	Port Usable	Port Down	Addr Change	Restart Failed	*Last Event Time
LCL	95219	12379748	0	95219	11237120	0	1*	0	0	0	14-FEB 11:34:31.70

```

EWA      455231  60696590      0   659096  167022916      0   1*   0   0   0   14-FEB 11:34:33.33
EIA      434778  57416246      0   606233  143615214      0   1*   0   0   0   14-FEB 11:34:41.63

```

MADDY PEA0 Device Errors 15-FEB-2009 11:43:12.89:

```

-- Transmit --
Generic      Hello
Device  Errors      Errors  Last Transmit Error
-----  -
LCL      0          0 (No time)
EWA      0          0 (No time)
EIA      0          0 (No time)

-----  Received -----
With Bad Mcast TR Short Short On Wrong Discard
Device  SYSTEMID Message CC Msg DX Msg Port PortDisa
-----  -
LCL      0          0          0          0          0          0
EWA      0          0          0          0          0          0
EIA      0          0          0          0          0          0

Allocation
Failures On
CH      VC
-----  -
LCL      0          0
EWA      0          0
EIA      0          0

```

MADDY LAN Counters and Errors 15-FEB-2009 11:43:12.89:

```

----- Transmit -----
Device  Packets      Bytes      Errors
-----  -
LCL      0          0          0
EWA     1734864     310115839  0
EIA     435056     65710688  0

----- Receive -----
Packets      Bytes      Errors
-----  -
LCL      0          0          0
EWA     2442265     480199487  0
EIA     903433     175593881  0

Buffer Unavailable
System      User
-----  -
LCL      0          0
EWA      0          0
EIA      0          0

-- Unrecognized --
Individual Multicast
-----  -
LCL      0          0
EWA      90       377506
EIA      0       297200

Events
-----  -
LCL      0
EWA      0
EIA      0

```

## Monitoring PEDRIVER for IP interfaces

The SDA command PE IP\_INTERFACE is useful for displaying PEDRIVER IP pseudo device data. Each IP device is a pseudo device on the system being used for NISCA communication.

SDA> PE IP\_INTERFACE

In the following example, the command displays the IEO interface's device summary.

### Example - SDA Command PE IP\_INTERFACE Display

SDA> PE IP\_INTERFACE

PE\$SDA Extension on I64MOZ (HP rx4640 (1.50GHz/6.0MB)) at 13-OCT-2008 15:10:08.03

I64MOZ Device Summary 13-OCT-2008 15:10:08.03:

```

Device Line Buffer MgtBuf Load Mgt Current Total Errors &
Device Type Speed Size SizeCap Class Priority IP Address Bytes Events Status
-----  -
IEO      0      1394      0      1000      0      16.138.182.6 264515808 48556 Run
Online XChain_Disabled

```

## Monitoring PEDRIVER Buses for IP interfaces

The SDA command SHOW PORT/BUS=BUS\_IP\_interface command is useful for displaying the PEDRIVER representation of an IP interface. To PEDRIVER, a bus is the logical representation of the IP interface. To list the names and addresses of buses, enter the SDA command SHOW PORT/ADDR=PE\_PDT and then press the Return key twice. The following example shows a display for the IP interface named IEO.

### Example - SDA Command SHOW PORT/BUS Display

SDA> show port/bus=886C0010

VMScluster data structures

```

--- BUS: 886C0010 (IE0) Device: IP IP Address: 10.138.182.6----- 1
Status: 00004203 run,online,xmt_chaining_disabled 2
----- Transmit ----- Receive ----- Structure Addresses ---
Msg Xmt      23459872773  Msg Rcv      24521301654  PORT Address      8850B9B8
Mcast Msgs      0      Mcast Msgs      0      VCIB Addr      886C02A0
Mcast Bytes      0      Mcast Bytes      0      HELLO Message Addr 886C02A0
Bytes Xmt      3055474713  Bytes Rcv      3545255112  BYE Message Addr 886C05CC

```

```

Outstand I/Os          0 Buffer Size          1394 Delete BUS Rtn Adr  90AA2EC8
Xmt Errors 5          0 Rcv Ring Size          0

--- Receive Errors --- ----- BUS Timer ----- ----- Datalink Events -----
TR Mcast Rcv          0 Handshake TMO    00000000 Last 22-SEP-2008 12:20:50.06
Rcv Bad SCSID         0 Listen TMO      00000000 Last Event          00004002
Rcv Short Msgs        0 HELLO timer     6 Port Usable       1
Fail CH Alloc         0 HELLO Xmt err   0 Port Unusable     0
Fail VC Alloc         0 Address Change   0
Wrong PORT            0 Port Restart Fail 0

```

Field	Description
1 IP Address	Displays the IP address of the interface.
2 Status	The Status line must always display the status as "online" to indicate that PEDRIVER can access its IP interface.
3 Msg Xmt (messages transmitted)	Shows the total number of packets transmitted over the virtual circuit to the remote node. It provides the Multicast (mcast) and the Multicast bytes that are transmitted.
4 Msg Rcv (messages received)	Shows the total number of packets received over the virtual circuit from the remote node. It provides the Multicast (mcast) and the Multicast bytes that are transmitted.
5 Xmt Errors (transmission errors)	Indicates the number of times PEDRIVER has been unable to transmit a packet using this IP interface.

### Monitoring PEDRIVER Channels for IP interfaces

The SDA command SHOW PORT/Channel=Channel\_IP\_interface command is useful for displaying the PEDRIVER representation of an IP interface. To PEDRIVER, a channel is the logical communication path between two IP interfaces located on different nodes. To list the names and addresses of channels created, enter the SDA command SHOW SYMBOL CH\_\* and then press the Return key twice. The following example shows a display for the IP interface named IE0.

### Example - SDA Command SHOW PORT/CHANNEL Display

```

SDA> show port/channel=CH_OOTY_IE0_WEO

VMScluster data structures
-----
-- PEDRIVER Channel (CH:886C5A40) for Virtual Circuit (VC:88161A80) OOTY --
State: 0004 open                      Status: 6F path,open,xchndis,rnhwavld,tight,fast
                                      ECS Status: Tight,Fast

BUS: 886BC010 (IE0) Lcl Device: IP     Lcl IP Address: 10.138.182.6 1
Rmt BUS Name: WEO Rmt Device: IP      Rmt IP Address: 11.146.235.10 2
Rmt Seq #: 0004 Open: 4-OCT-2008 00:18:58.94 Close: 4-OCT-2008 00:18:24.53

- Transmit Counters --- - Receive Counters ---- - Channel Characteristics --
Bytes Xmt      745486312 Bytes Rcv      2638847244 Protocol Version      1.6.0
Msg Xmt        63803681 Msg Rcv        126279729 Supported Services    00000000
  Ctrl Msgs          569   Ctrl Msgs          565 Local CH Sequence #    0003
  Ctrl Bytes        63220   Ctrl Bytes        62804 Average RTT (usec)    5780.8
                               Mcast Msgs        106871 Buffer Size:
                               Mcast Bytes    11114584   Current              1394
- Errors ----- Remote              1394
Listen TMO          2 Short CC Msgs      0 Local              1394
TR ReXmt            605 Incompat Chan     0 Negotiated         1394
DL Xmt Errors       0 No MSCP Srvr     0 Priority            0
CC HS TMO           0 Disk Not Srvd   0 Hops                2
Bad Authorize       0 Old Rmt Seq#    0 Load Class         100
Bad ECO             0 Rmt TR Rcv Cache Size 64
Bad Multicast       0 Rmt DL Rcv Buffers   8
                               Losses              0
- Miscellaneous ----- - Buf Size Probing----- - Delay Probing -----
Prv Lstn Timer      5 SP Schd Timeout   6 DP Schd Timeouts   0

```

```

Next ECS Chan  886C5A40  SP Starts          1  DP Starts          0
                  SP Complete        1  DP Complete        0
- Management -----  SP HS TMO          0  DP HS TMO          1
Mgt Priority    0        HS Remaining Retries 4
Mgt Hops       0        Last Probe Size    1395
Mgt Max Buf Siz 8110

```

Field	Description
1 Lcl IP Address (Local IP Address)	Displays the IP address of the local interface.
2 Rmt IP Address (Remote IP Address)-	Displays the IP address of the remote interface.

## Appendix B Restrictions and Known Problems in this Release

This section provides the restrictions and known problems with Cluster over IP in this release.

- Configuring a Node While Upgradation**  
 During upgradation from the prior versions to OpenVMS 8.4, Cluster over IP cannot be enabled. OpenVMS, while upgradation does not call CLUSTER\_CONFIG[\_LAN], which is required for enabling Cluster over IP. Hence, the node joins the existing cluster of which it was the member before upgrade.  
 For enabling Cluster over IP, CLUSTER\_CONFIG\_LAN procedure must be called explicitly after up gradation.  
 This restriction will be fixed in the future release.
- Alpha Satellite Node Requires LAN channels With Disk Server**  
 Alpha satellite boot fails in an IP only environment. That is, while booting an Alpha satellite, if all the nodes, including the boot servers, are using only IP channels for cluster communication, the satellite boot fails with the following message:

```

cluster-W-PROTOCOL_TIMEOUT, NISCA protocol timeout %VMScluster-I-REINIT_WAIT,
Waiting for access to the system disk server

```
- IPv6 Support**  
 Cluster over IP does not support IPv6 type address for cluster communication interface.
- Dynamic Host Configuration Protocol (DHCP) or Secondary Address Support**  
 Cluster over IP requires the addresses that are used for cluster communication, which are static, primary address on that interface. Furthermore, IP address used for cluster communication must not be used for Failsafe configuration.
- Multiple IP Interface Configuration**  
 If you configure multiple IP interface with the same default gateway, loss of communication on any interface may result in disrupted cluster communication with CLUEXITS.
- Integrity servers Satellite Node and Bootserver in the Same LAN**  
 An Integrity server satellite node must be in the same LAN as its boot server for the satellite node to initialize cluster over IP successfully and to join the cluster successfully.

It is also necessary to have LAN cluster communication between Integrity server satellite node and the boot server for the satellite node to be able to initialize cluster over IP during the satellite bootup.

- Requires OpenVMS Version 8.4 and HP TCP/IP Services Version 5.7  
Cluster over IP is available only on OpenVMS Version 8.4 Alpha and Integrity servers. Cluster over IP also requires HP TCP/IP services for OpenVMS, Version 5.7.
- Reactive IP Interface Results in Loss of Connection  
If the interface used for cluster communication is reactivated by ifconfig, it results in losing cluster communication to other nodes and also results in cluexit of nodes.
- Multiple Gateway Configuration  
Cluster over IP configuration information is stored in the configuration files, which are loaded early in the boot time. This configuration information also includes the default route or gateway that is used by TCP/IP. Currently, only one default route can be entered in the configuration file and used during the node bootup.
- Block Transfer XMIT Chaining  
PEdriver emulates each IP interface used for cluster communication with a BUS. An IP bus will have the characteristics of Xchain\_Disabled status as shown. The means that the block transfer packets transmitted through TCP/IP are copied from the PEdriver to the TCP/IP buffers.

```
$ MC SCACP SHOW IP
NODEG PEA0 Device Summary 16-FEB-2009 12:29:15.92:
```

DeviceType	Errors + Events	Status	Mgt Priority	Buffer Size	MgtMax BufSiz	Line Speed	Total Pkts(S+R)	Current IP Address
IE0	184	Run Online	XChain_Disabled	0	1394	0	N/A	1419711 15.146.235.222

- LANCP for Downline Load  
Cluster over IP requires LANCP, instead of DECnet for downline load since the changes related to configuring cluster over IP and enabling cluster over IP is available only with CLUSTER\_CONFIG\_LAN.COM. This restriction will be fixed in the future releases of HP Clusters.
- Duplex Mismatch  
A duplex mode mismatch or a change in duplex mode from half to full on the host duplex can result in CLUEXIT when IP is used for cluster communication. It is recommended to check for the duplex mismatch issues to avoid cluexit.

## For more information

For more information about the OpenVMS Cluster Systems, see *HP OpenVMS Cluster Systems Guide*.

For the cluster guidelines and recommendations, see *Guidelines for OpenVMS Cluster Configuration Manual*.

For more information about SCACP commands, see *HP OpenVMS System Management Utilities Reference Manual*.

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

November 2009