HP NonStop Networking Overview

HP Part Number: 529874-009
Published: August 2012
Edition: J06.04 and all subsequent J-series RVUs, H06.16 and all subsequent H-series RVUs

Legal Notice

Contents

Figures

Tables

About This Guide

This manual provides an overview of networking on the J-series and H-series release version updates (RVUs) on Integrity NonStop systems, including supported hardware configurations and migration and planning information.

Supported Release Version Updates (RVUs)

This manual supports J06.04 and all subsequent J-series RVUs and H06.16 and all subsequent H-series RVUs until otherwise indicated in a replacement publication.

Intended Audience

This manual is for network administrators new to the NonStop system environment and for those migrating from the G-series RVUs to the H-series or J-series RVUs, information technology planners (both customer and sales support), and Global NonStop Solution Center (GNSC) analysts looking for information about the Integrity NonStop systems.

For information about other kinds of networking solutions, the Expand subsystem, and background information about NonStop system concepts and components, see the *Introduction to Networking for HP NonStop S-Series Servers*.

New and Changed Information in This Edition

This edition adds support for:

- Network partitioning through use of multiple PROVIDERs on CLIMs
- NonStop NS2100 systems

Changes in this edition are marked with change bars.

New and Changed Information in Previous Editions

The previous edition added support for the IB CLIM and the NonStop NS2200-series systems.

The previous edition added support for token ring on HP Integrity NonStop BladeSystem and incorporated several corrections and enhancements.

Previous editions included updates for the HP NonStop BladeCluster Solution. See

- “Related Information” (page 6)
- “Networking Hardware Products Available on H-Series and J-Series RVUs” (page 7)
- “Platform Interoperability ” (page 11)

Information about using the NonStop S-series I/O enclosure for SS7 functionality on Integrity NonStop BladeSystems was added to “Migrating From a Platform Other Than a NonStop Server” (page 33).

Document Organization

Title	Content
“Networking on Integrity NonStop Systems” (page 7)	Describes fundamentals of networking on Integrity NonStop Systems. Provides network interoperability topology diagrams.
“Networking Concepts ” (page 24)	Describes IP networking concepts.
“Planning for Migrating Networking Solutions to H-Series and J-Series RVUs” (page 29)	Describes migrating your networking implementations from a G-series RVUs to H-series and J-series RVUs as well as from other environments.

Related Information

To plan for networking on Integrity NonStop systems, use this manual. Depending on the tasks you are performing, you might also need these related manuals:

- *AWAN 3886 Server Installation and Support Guide*
- *CLuster I/O Module (CLIM) Software Compatibility Reference*
- *BladeCluster Solution Manual*
- *Cluster I/O Protocols (CIP) Configuration and Management Manual*
- *DNS Configuration and Management Manual*
- *LAN Configuration and Management Manual*
- *Expand Configuration and Management Manual*
- *Gigabit Ethernet 4-Port Adapter Installation and Support Guide*
- *Introduction to Networking for NonStop S-Series Servers*
- *SWAN Concentrator Installation and Support Guide*
- *SWAN 2 Concentrator Installation and Support Guide*
- *TCP/IP Configuration and Management Manual*
- *TCP/IPv6 Configuration and Management Manual*
- *WAN Subsystem Configuration and Management Manual*

Also see the planning guide for your system (for example, the *NonStop BladeSystem Planning Guide)* for a list of all related manuals.

All of these manuals are available at the HP Business Support Center (http://www.hp.com/go/nonstop-docs).

Publishing History

Part Number	Product Version	Publication Date
529874-006	N.A.	August 2009
529874-007	N.A.	March 2011
529874-008	N.A.	August 2011
529874-009	N.A.	August 2012

HP Encourages Your Comments

HP encourages your comments concerning this document. We are committed to providing documentation that meets your needs. Send any errors found, suggestions for improvement, or compliments to:

docsfeedback@hp.com

Include the document title, part number, and any comment, error found, or suggestion for improvement you have concerning this document.

1 Networking on Integrity NonStop Systems

This section provides a high-level overview of networking on H-series and J-series RVUs on Integrity NonStop systems. Topics discussed in this section include:

- Networking Hardware Products Available on H-Series and J-Series RVUs
- Networking Software Products Available on H-Series and J-Series RVUs (page 10)
- Platform Interoperability (page 11)
- What's Unique About Networking on Integrity NonStop Systems (page 14)
- What's Unique About Networking on Integrity NonStop BladeSystems (page 21)
- Summary Comparison of NonStop TCP/IP Products on H-Series and J-Series RVUs (page 21)

Networking Hardware Products Available on H-Series and J-Series RVUs

NOTE: The installation slots for G4SAs, FCSAs and CLIMs differ by system model. For more information, see the planning guide for your system model (for example, the *NonStop NS16000 Series Planning Guide*).

Products for Ethernet Connectivity

Ethernet connectivity on H-Series and J-Series is provided by either:

- A Gigabit Ethernet 4-port ServerNet adapter (G4SA) installed in an IOAM enclosure of a NonStop NS-series system or NonStop BladeSystem.
- A networking CLuster I/O (CLIM): IP or Telco CLIM installed in or connected to a NonStop NS-Series or NonStop BladeSystem

The G4SA is the only Ethernet adapter that can be installed in an I/O Adapter Module (IOAM) enclosure. There are two IOAMs in an IOAM enclosure.

The Networking CLIM provides Ethernet connectivity for these H-series and J-series environments and systems:

- H-series RVUs on NonStop NS16000 series systems and J-series RVUs on NonStop BladeSystems, where a Networking CLIM can be used instead of a G4SA and, when used in combination with a Storage CLIM, replaces all IOAM functionality.
- J-series RVUs on NonStop NS2000, NS2100, and NS2200 series systems, where it can be used in combination with the Versatile I/O (VIO). (In this environment, the Storage CLIM is required because all storage is CLIM-attached. See the *NonStop Storage Overview* for more information about storage options.)

Products for InfiniBand Connectivity

InfiniBand connectivity is provided by the IB CLIM and is supported on NonStop BladeSystems NB50000c and NB54000c for J06.12 and later J-series RVUs and NonStop NS16000 systems for H06.23 and later H-series RVUs.

NOTE: IB CLIMs are only used as a Low Latency Solution. They do not provide general purpose InfiniBand connectivity for NonStop Systems.

Products for Connecting to NonStop I/O Enclosures

You can use other networking hardware products by connecting some Integrity NonStop system models to a NonStop S-series I/O enclosure that has IOMF2 components installed. This connection is made from the system processor switch (p-switch) or ServerNet switch. This option is not available on the NonStop NS2000 series system or NonStop NS2100 system.

Figure 1 Integrity NonStop System P-Switch Connection to a NonStop S-Series I/O Enclosure

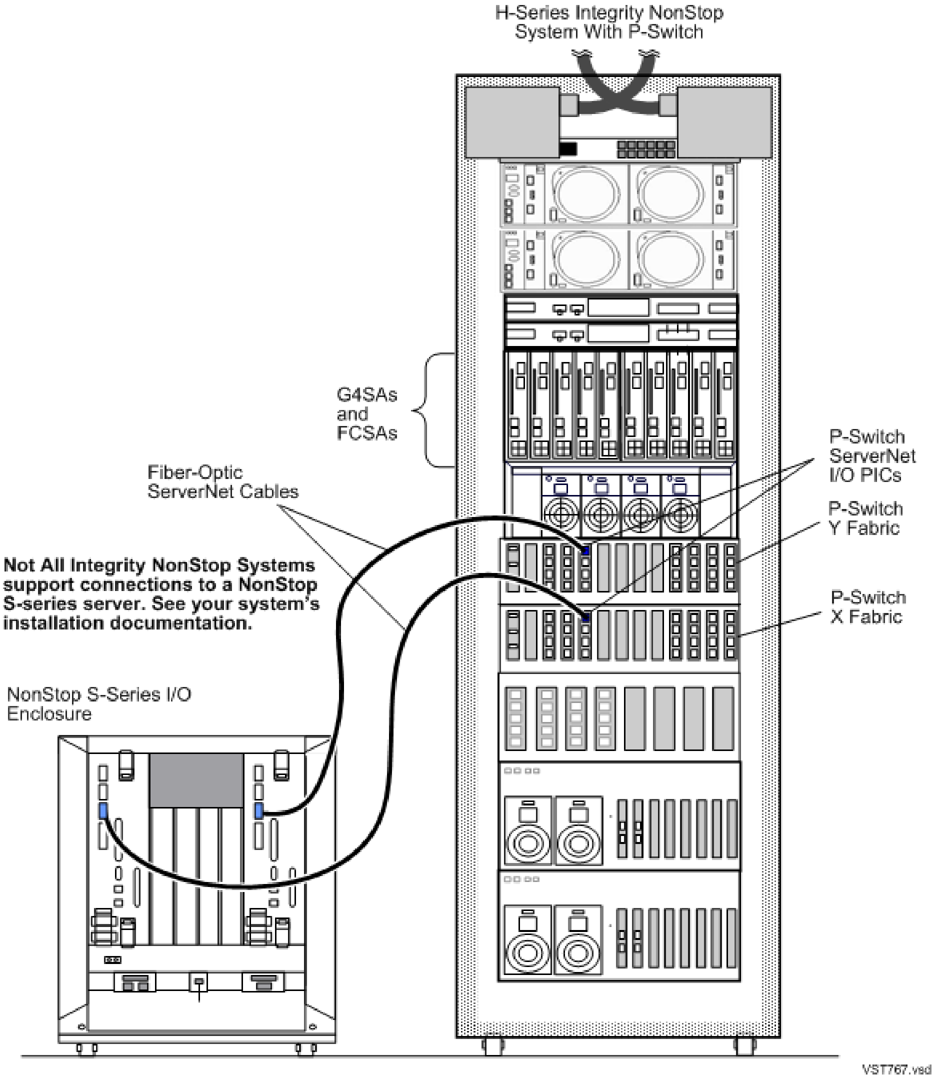
H-Series Integrity NonStop System With P-Switch
G4SAs and FCSAs
Fiber-Optic ServerNet Cables
Not All Integrity NonStop Systems support connections to a NonStop S-series server. See your system's installation documentation.
NonStop S-Series I/O Enclosure
P-Switch ServerNet I/O PICs
P-Switch Y Fabric
P-Switch X Fabric
VST767.vsd

Figure 2 Integrity NonStop BladeSystem Connection to a NonStop S-Series I/O Enclosure

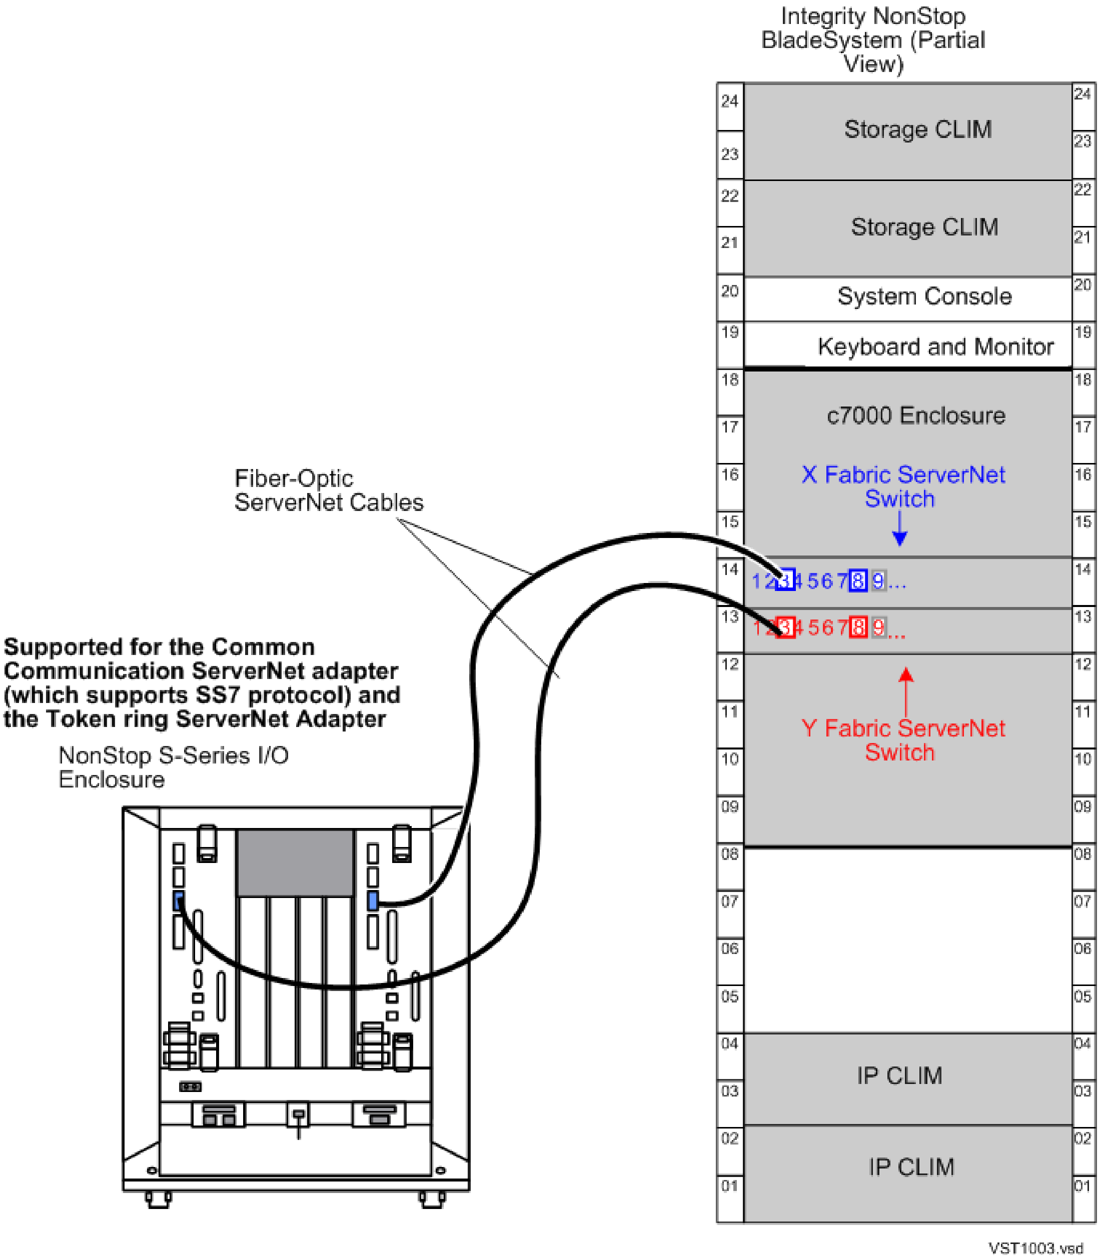
Integrity NonStop BladeSystem (Partial View)
Storage CLIM
Storage CLIM
System Console
Keyboard and Monitor
c7000 Enclosure
X Fabric ServerNet Switch
1 2 3 4 5 6 7 8 9...
1 2 3 4 5 6 7 8 9...
Y Fabric ServerNet Switch
IP CLIM
IP CLIM
24
23
22
21
20
19
18
17
16
15
14
13
12
11
10
09
08
07
06
05
04
03
02
01
Fiber-Optic ServerNet Cables
Supported for the Common Communication ServerNet adapter (which supports SS7 protocol) and the Token ring ServerNet Adapter
NonStop S-Series I/O Enclosure
VST1003.vsd

Legacy NonStop S-series networking I/O adapters supported through the NonStop S-series I/O enclosure include:

- GESA, FESA, and E4SA (Ethernet adapters)
- TRSA (Token-ring)
- ATM3SA (Asynchronous Transfer Mode 3–port ServerNet adapter)

NOTE: Support for these products is only for legacy adapters already existing in the environment. These products have reached end-of-sale status.

NOTE: The only supported S-series I/O connections for the NonStop BladeSystems are CCSA and TRSA.

The Versatile I/O (VIO) enclosure provides connectivity on some systems. The VIO enclosure supports Ethernet and, for some systems, fibre channel connectivity.

On NonStop NS-series systems, ServerNet wide area network (SWAN) concentrators use TCP/IP to connect through the G4SA or CLIM, through Ethernet ports in the VIO enclosure or through other Ethernet adapters in a NonStop S-series I/O enclosure. On NonStop BladeSystems, SWAN concentrators connect through the CLIM or G4SA. On NonStop NS2000, NS2100, and NS2200-series systems, SWAN concentrators connect through the VIO. AWAN is not supported on NonStop NS2000, NS2100, and NonStop NS2200-series systems.

Products that Interconnect NonStop Systems — BladeCluster Solution

The BladeCluster Solution is comprised of network topologies that interconnect NonStop BladeSystems and NonStop NS16000 series systems as nodes, which can also cluster with 6770/6780 ServerNet clusters. These interconnected nodes use ServerNet to pass information from node–to–node functioning as one large processing entity. The BladeCluster Solution offers five network topologies:

- BladeCluster/BladeSystem Solution (all BladeSystems)
- BladeCluster/NonStop NS16000 Series Solution (all Non Stop NS16000 series systems)
- BladeCluster/Mixed System Solution (NonStop BladeSystems and NS16000 series systems)
- BladeCluster/6700 Series Solution (BladeCluster clustered with 6770/6780 ServerNet Cluster)
- BladeCluster/Multi-Zone Solution (Up to 3 zones and up to 24 nodes – short-haul and long-haul distance options)

To support the BladeCluster Solution topologies, each BladeCluster node is connected along each of two ServerNet fabrics via 2-way multilane ServerNet 3 links (unless configured for ServerNet 2 in interzone distances greater than 25 km).

The Multi-Zone long-haul topology offers distance options of either 25 km, 50 km, or 65 km.

Hardware requirements for the BladeCluster Solution include specific switches for BladeCluster (ServerNet switch for BladeCluster and Processor Switch for BladeCluster) and 1 Advanced Cluster Hub (ACH) per fabric. The ACH is used to connect the X and Y fabric processor switches and ServerNet switches in a BladeCluster topology. For the multi-zone long-haul topology, Dense Wavelength Division Multiplexer (DWDM) third-party equipment is required.

The DWDM routes the ServerNet lanes to their appropriate zone and one to two DWDMs are required per zone. Contact your HP reprensentative for a list of DWDMs that interoperate with the BladeCluster Solution.

For detailed BladeCluster software and licensing requirements, refer to the *BladeCluster Solution Manual*. The manual also includes detailed installation, connection, migration, and upgrade procedures for all BladeCluster topologies.

Networking Software Products Available on H-Series and J-Series RVUs

Most of the networking software available for the G-series RVUs is also available on H-series and J-series RVUs, including:

- AM3270 Access Method
- Asynchronous Terminals and Printer Processes
- ATP6100 WANPRINT
- Cluster I/O Protocols (CIP) (Not all Integrity NonStop systems support CIP; see “Networking Hardware Products Available on H-Series and J-Series RVUs” (page 7))
- DNS 9.x
- Envoy
- EnvoyACP/XF
- Expand

- NonStop TCP/IPv6
- NonStop TCP/IP (conventional)
- OSI/AS
- OSI/FTAM
- OSI/MHS
- OSI/TS
- Port Access Method (PAM)
- QIO
- ServerNet LAN Systems Access (SLSA)
- SNAX High Level Support (SNAX/HLS)
- SNAX/XF (includes SNAX/APN functionality)
- SNMP
- Spooler
- Spooler FastP Network Print Processes
- Spooler Plus
- TELSERV
- TN3270e
- TR3271 Tributary Access Method
- Wide area network (WAN) subsystem

Networking software available only on G-series RVUs includes:

- 6100 TINET Multi-PT Supervisor-PDG
- 6100 UTS-40 Multi-PT Supervisor
- 6100 UTS-40 Multi-PT Tributary
- 6100 VIP Multi-PT Supervisor
- Domain Name Server (DNS) (T6021) (Replaced by DNS 9.x)
- Enform
- Fiber Optic Extension (FOX) Gateway
- Multilan Access method and File Server
- NonStop IPX/SPX
- Novell LAN Print Spooler
- Parallel Library TCP/IP (functionality available in NonStop TCP/IPv6 on Integrity NonStop systems)
- TANDEM NBT NETBIOS for IPX
- TANDEM NBT NETBIOS for TCP

Platform Interoperability

Integrity NonStop systems can communicate directly with a NonStop K-series system over Expand. Some Integrity NonStop NS-series systems can also communicate indirectly with a NonStop K-series system through a NonStop S-series system or a ServerNet cluster switch.

NOTE: Not all Integrity NonStop systems support ServerNet clustering. For more information, see the planning guide for your system model (for example, the *Integrity NonStop NS16000 Series Planning Guide*).

Figure 3 shows interoperability between multiple NonStop systems.

Figure 3 Connecting Legacy Systems With Integrity NonStop Systems Running H-Series RVUs

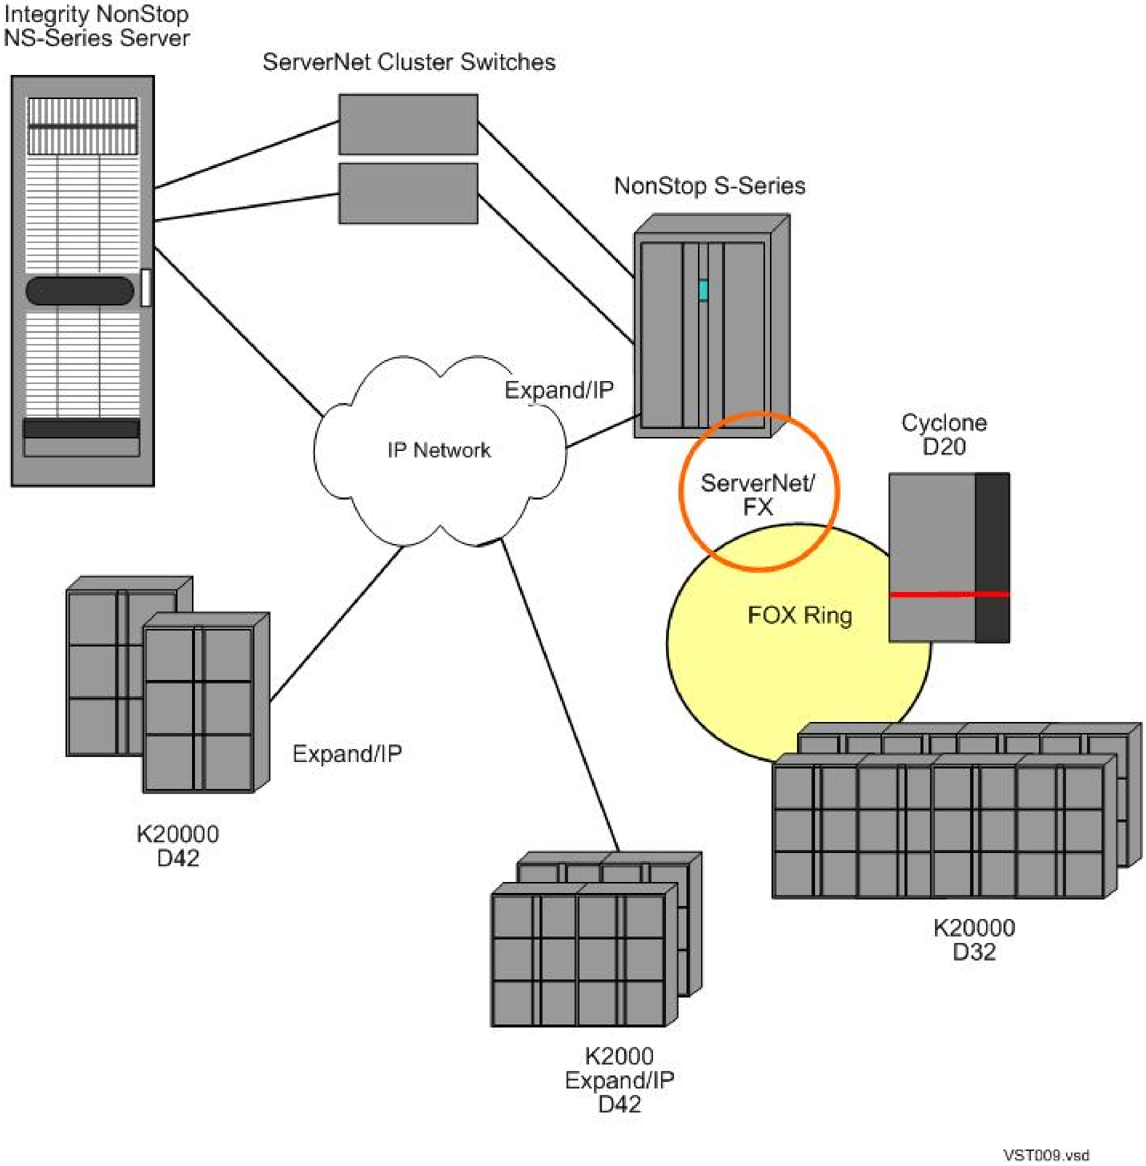

The ServerNet cluster switch connects to the p-switch of some Integrity NonStop systems and can interoperate with nodes in the BladeCluster Solution Release 1.1.

For information about including the Integrity NonStop system in a ServerNet cluster, see the *ServerNet Cluster Supplement for NonStop NS-Series Servers*. For information about including Integrity NonStop systems in a BladeCluster Solution, see the *BladeCluster Solution Manual*.

NonStop Integrity Systems can support these clustering topologies:

Cluster Topology	Systems Supported	Hardware Required
BladeCluster Solution	NonStop BladeSystems	Refer to the *BladeCluster Solution Manual*.
	Mixed NonStop BladeSystems and NS16000 series systems	Refer to the *BladeCluster Solution Manual*.
	Connection to 6780 ServerNet Cluster comprised of NonStop BladeSystems, NS16000 series systems, NS14000 series systems, or NonStop S-series systems	Refer to the *BladeCluster Solution Manual*.
	Connection to 6770 ServerNet Cluster comprised of NonStop BladeSystems, NS16000 series systems, NS14000 series systems, or NonStop S-series systems	Refer to the *BladeCluster Solution Manual*.
6780 ServerNet cluster	NonStop BladeSystems NonStop NS16000 NS14000 series systems Connection to BladeCluster Solution comprised of NonStop BladeSystems and NS16000 series systems Connection to 6770 ServerNet Cluster comprised of NonStop BladeSystems, NS16000 series systems, NS14000 series systems, and NonStop S-series systems	6780 ServerNet cluster switch for all these topologies
6770 ServerNet cluster	NonStop BladeSystems NonStop NS16000 series systems NS14000 series systems Connection to BladeCluster Solution comprised of NonStop BladeSystems and NS16000 series systems Connection to 6780 ServerNet Cluster comprised of NonStop BladeSystems, NS16000 series systems, NS14000 series systems, and NonStop S-series systems	6770 ServerNet cluster switch for all these topologies

Figure 4 shows interoperability between multiple NonStop systems but does not show the BladeCluster Solution. For figures showing the BladeCluster Solution, see the *BladeCluster Solution Manual*.

Figure 4 Connecting Legacy Systems With Integrity NonStop BladeSystems

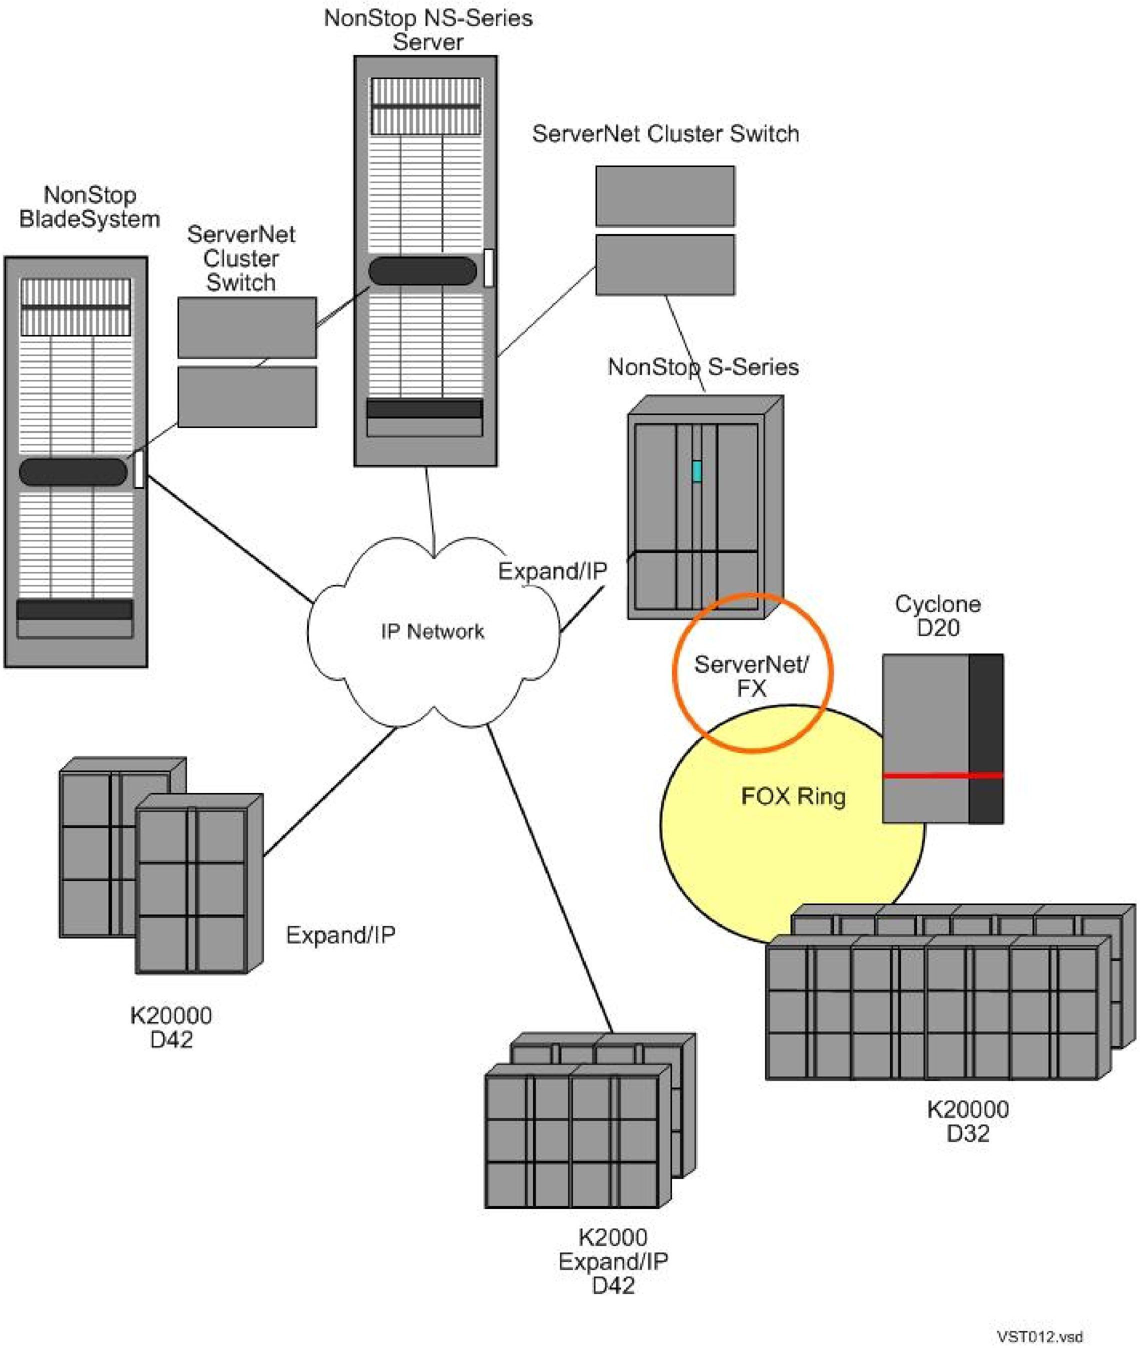

What's Unique About Networking on Integrity NonStop Systems

If you are new to the Integrity NonStop systems, you should be aware of differences between its networking architecture and solutions and the architecture of other platforms, including other NonStop systems. This subsection explains:

- Integrity NonStop System Networking Compared to NonStop S-Series Server Networking
- Integrity NonStop System Networking Compared to Other Platforms

Integrity NonStop System Networking Compared to NonStop S-Series Server Networking

The main difference to be aware of between the NonStop S-series system and the Integrity NonStop systems is that some of the networking technology of the NonStop S-series system has been ported to the Integrity NonStop system, but the hardware for that legacy networking technology is accessible only by connecting to a NonStop S-series I/O enclosure. (See Networking Software Products Available on H-Series and J-Series RVUs (page 10) for a list of all networking products available.)

Integrity NonStop System Networking Compared to Other Platforms

The Integrity NonStop system clustering technology makes it possible for the system to appear as multiple hosts. Also, the system, when combined with the parallel processing architecture of the NonStop TCP/IPv6 or CIP product, can appear as a single host while allowing all 16 processors to service one port. The latter feature allows for high scalability of TCP/IP client applications.

The NonStop TCP/IPv6 and CIP products offer the unique ability to bind multiple servers to a single port. The listening application instances can run in different processors, bind to a shared IP address and port, and so service a single IP image with 16 processors of computing power.

This feature, called round-robin filtering, allows for scalability of the listening application (in this example, the web server process). By sharing a port, the web server process can provide a single-IP host image to the world through either a single IP address or, in the multi-homed situation, using DNS round-robin address rotation through multiple IP addresses. In the scenario where the web server processes are sharing an IP address, the answer that the authoritative name server process provides for the web server process is the same for all 16 instances of that application server. Requests that come into the web server process at that IP address are distributed among the 16 web server processes.

For example, you can run a website using iTP Secure WebServer on your Integrity NonStop system, enable round-robin filtering with NonStop TCP/IPv6 or CIP, and benefit in two ways from the parallel TCP/IP architecture. First, you can run multiple httpd server processes (the listening processes) each bound to the same port and IP address; this action scales your httpd server process up to 16 processors. Second, no interprocess hop occurs from the processor containing the listening process and a processor with TCP/IP access to the communication adapter. CIP extends this functionality to allow more than one process per processor to serve a single IP address.

NOTE: Round-robin filtering needs to be enabled for NonStop TCP/IPv6 and CIP.

NOTE: NonStop TCP/IPv6 had a limitation of one listening process per processor per port. CIP does not have this limitation. There can be many listening processes per processor per port.

Additionally, the NonStop TCP/IPv6 and CIP architectures allow direct access from each processor in a node to the adapter or CLIM. Direct access to the adapter is distinct from the architecture of conventional NonStop TCP/IP, which can involve an interprocess hop from a processor containing the TCP/IP stack to the processor containing the application.

For more information about the parallel TCP/IP architecture available on NonStop systems, about the models for listening applications and how they benefit from this architecture, and for procedures for configuring the networking environment for round-robin filtering, see the *TCP/IPv6 Configuration and Management Manual*. For information about configuring CIP, see the *Cluster I/O Protocols (CIP) Configuration and Management Manual*.

NOTE: The models for listening applications described in the *TCP/IPv6 Configuration and Management Manual* also apply to CIP.

For information about the iTP Secure WebServer, see the *iTP Secure WebServer System Administrator's Guide*.

The remainder of this subsection details the unique characteristics of networking in the NonStop system environment.

You can use various methods for scaling your DNS implementation, some of which are unique to the NonStop system and some of which are enhanced by the NonStop system architecture. You can also use various methods for scaling your network interface capacity as well as scaling your application through DNS capabilities.

For the purposes of this discussion we make the distinction between:

- Network Scalability
- Application Scalability

Network Scalability

Network scalability refers to the use of multiple physical interfaces to accommodate bandwidth requirements for which a single network interface is insufficient. You can achieve network scalability by having multiple network interfaces on multiple hosts or by having a multi-homed host (a single system that has multiple network interfaces). You can make this form of scalability either explicit or transparent. If explicit, then both ends of the connection simply accommodate multiple interfaces by specifying a list of IP addresses.

To make network scalability transparent with a multi-homed host or with multiple hosts, you can use the Domain Name System (DNS) and have the multiple interfaces share a single, externally visible name. One way to have multiple interfaces share a name, is to have a DNS name server process on each of the subnets and configure each of those name server processes to return the IP address of the host that is also on that subnet (Figure 5).

Figure 5 Transparent Network Scalability, Multiple Hosts

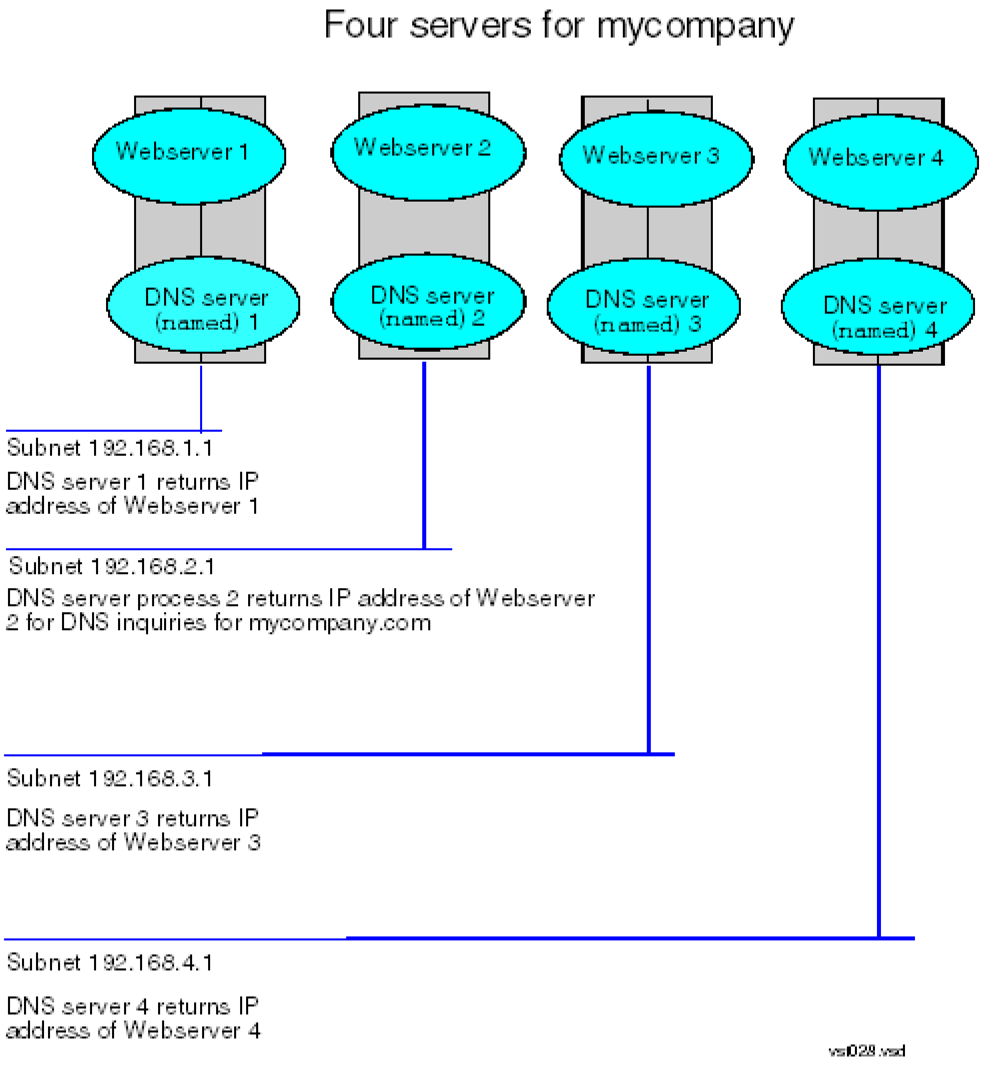
Four servers for mycompany
Webserver 1
Webserver 2
Webserver 3
Webserver 4
DNS server (named) 1
DNS server (named) 2
DNS server (named) 3
DNS server (named) 4
Subnet 192.168.1.1
DNS server 1 returns IP address of Webserver 1
Subnet 192.168.2.1
DNS server process 2 returns IP address of Webserver 2 for DNS inquiries for mycompany.com
Subnet 192.168.3.1
DNS server 3 returns IP address of Webserver 3
Subnet 192.168.4.1
DNS server 4 returns IP address of Webserver 4
vst028.vsd

In the example shown in Figure 5, there are four DNS name server processes, one on each subnet. Each name server process is configured to return the IP address of the host's interface on that subnet. All clients requesting an IP address on that subnet are returned the IP address of the host that is also on that subnet.

- Clients on subnet 192.168.1.1 requesting the IP address of MyCompany.com receive the IP address of Webserver 1.
- Clients on subnet 192.168.2.1 requesting the IP address of MyCompany.com receive the IP address of Webserver 2.
- Clients on subnet 192.168.3.1 requesting the IP address of MyCompany.com receive the IP address of Webserver 3.
- Clients on subnet 192.168.4.1 requesting the IP address of MyCompany.com receive the IP address of Webserver 4.

Figure 5 (page 16) shows transparent network scalability on multiple hosts.

The other way of achieving transparent network-level scaling is to use DNS round-robin address rotation where one DNS name server process returns different IP addresses for a given resource service. With this technique, you would ensure there is a different network interface associated with each IP address; one name server process would then answer requests for the name with those IP addresses on a rotating basis. In Figure 6, the DNS server process rotates Webserver requests for Mycompany.com among interfaces 1.1.1.1, 1.1.2.1, 1.1.3.1 and 1.1.4.1. (On a multi-homed NonStop server, the separate systems shown inFigure 6 are contained in one system.)

Figure 6 DNS Round-Robin Address Rotation

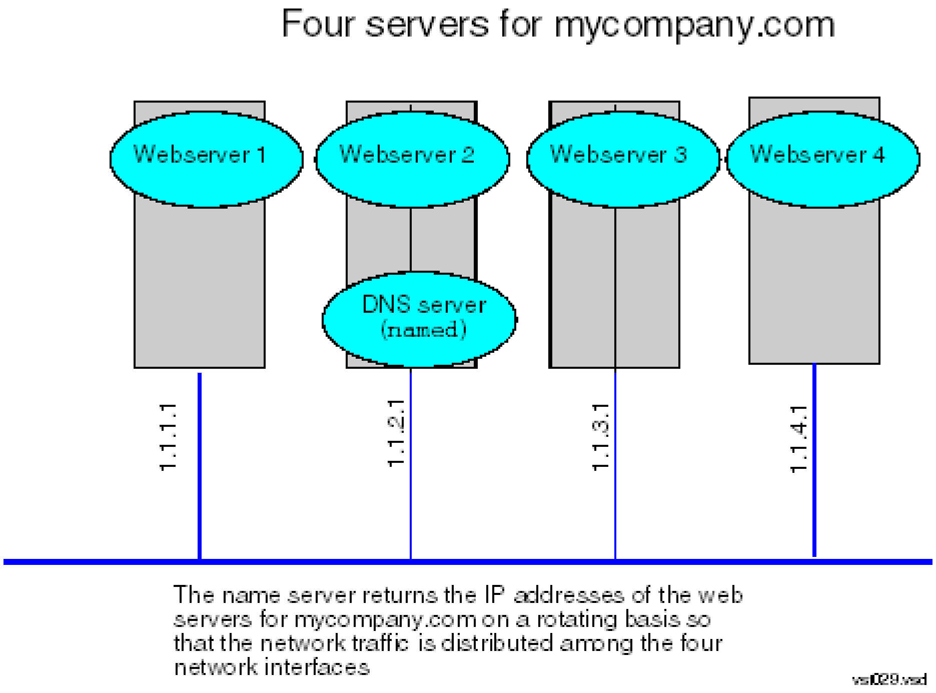

The name server returns the IP addresses of the web servers for mycompany.com on a rotating basis so that the network traffic is distributed among the four network interfaces

See the *DNS Configuration and Management Manual* for more information about configuring a DNS name server process on your NonStop system.

Application Scalability

Application scalability refers to distributing application-service load across multiple instances of the application server process. In the example of the scaling mechanism provided by DNS, the service application (the web server process) is limited when the server instances need to share data that is not resident on all the hosts. For example, a set of web server processes distributed over n hosts must either have the data it is servicing duplicated across those hosts or connect to a back-end shared database so that all the web server processes are accessing up-to-date information.

In a NonStop system environment, those n hosts can be consolidated onto a single system and have access to a shared database on that same system.

Returning to the Nonstop TCP/IPv6 and CIP subsystems' architectural features of allowing multiple server processes to bind to a shared port and IP address, the web server process in can run 16 copies of itself and access a shared database while presenting a single image to the network.

NOTE: NonStop TCP/IPv6 had a limitation of one listening process per processor per port. CIP does not have this limitation. There can be many listening processes per processor per port.

Figure 7 shows round-robin filtering on a single IP address. Incoming connection requests are distributed among 16 instances of the server process in 16 different processors.

Figure 7 Round-Robin Filtering in the NonStop TCP/IPv6 and CIP Environments

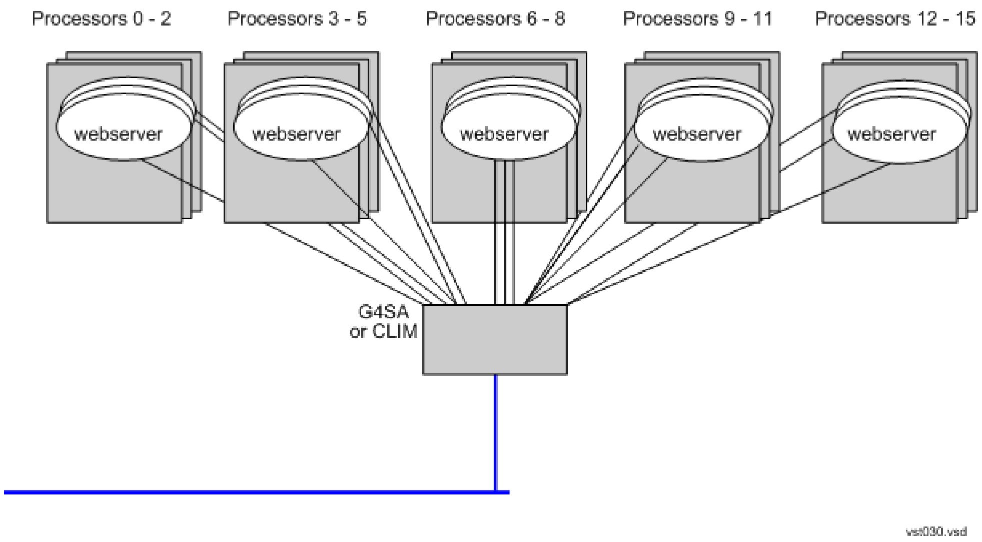

Scaling Your Application on the NonStop System

Several listener models are discussed in the *TCP/IPv6 Configuration and Management Manual*, each with unique configuration requirements depending on how that particular listener model works. CIP also supports the listening models and scaling techniques described in the *TCP/IPv6 Configuration and Management Manual*. By configuring your NonStop TCP/IPv6 or CIP subsystem to use round-robin filtering, you can cause incoming connection requests to be distributed among different application instances running in different processors.

NOTE: To enable NonStop TCP/IPv6 and CIP round-robin filtering only for certain applications, you can limit the ports upon which round-robin filtering is enabled by adding a DEFINE. See the *TCP/IPv6 Configuration and Management Manual* or the *Cluster I/O Protocols (CIP) Configuration and Management Manual* for an explanation of this DEFINE.

Round-robin filtering configuration for an application requires that you configure the NonStop TCP/IPv6 or CIP environment with round-robin filtering and that you define the transport-service providers (TCP6SAM or CIPSAM processes) in the same TACL session in which you define the filter key.

You configure NonStop TCP/IPv6 or CIP with round-robin filtering enabled in the Guardian environment and the OSS environment inherits the DEFINE that establishes the environment.

Fault-Tolerance and Scalability on the NonStop System

When you have configured for scalability, the failure of one or more processors does not impact the availability in any way:

- Availability of an application is not impacted by a failed processor because the application process is replicated.
- Availability of Internet accessibility is not impacted by a failed processor because you have running IP stacks configured on the running processors with the IP address still available so clients do not have to talk to an alternate IP address.
- Recovery of scalability is not affected by a failed processor because when the processors are reloaded and applications are restarted, the applications automatically resume their scaled configuration without disruption to the single IP view. The re-instantiated applications start using the IP address that had been configured previously.

Failover: Fault-Tolerant Connectivity

Failover allows a second interface to take over connections from an interface in a pair configured for failover if the other interface fails.

Unlike other failover implementations, NonStop TCP/IPv6 and CIP failover do not require one of the interfaces to act as a “hot standby” in anticipation of a failure. Both interfaces are active and allow inbound and outbound network traffic to be distributed between them. Hence, in addition to having fault-tolerance at the interface level by using failover, you also gain scalability when all your interfaces are performing correctly.

For more information about failover, see the *TCP/IPv6 Configuration and Management Manual* or the *Cluster I/O Protocols (CIP) Configuration and Management Manual.*.

The CIP subsystem can be configured to provide failover between interfaces on a CLIM and can provide failover between two CLIMs. Failover between two CLIMs results in a loss of TCP and SCTP connections so applications need to reestablish those connections. However once sessions have been reestablished, all other configuration aspects of the failed-over interface, are restored.

Logical Network Partitioning (LNP) and Providers

In the NonStop TCP/IPv6 environment, all applications have access to all interfaces (IP addresses) unless you configure logical-network partitioning (LNP). The CIP environment shares this feature, except you use Providers to restrict applications.

One of the most important differences between conventional TCP/IP and NonStop TCP/IPv6 and CIP is that NonStop TCP/IPv6 and CIP have one manager process ($ZZTCP or $ZZCIP) and all interfaces are associated with that single process. This difference influences how you configure the subsystem. In conventional TCP/IP, you can have multiple TCP/IP processes, each having one or more interfaces uniquely associated with them, as shown in Figure 8.

Figure 8 Conventional TCP/IP: Data From the Interface is Restricted to Applications Using the Associated Process

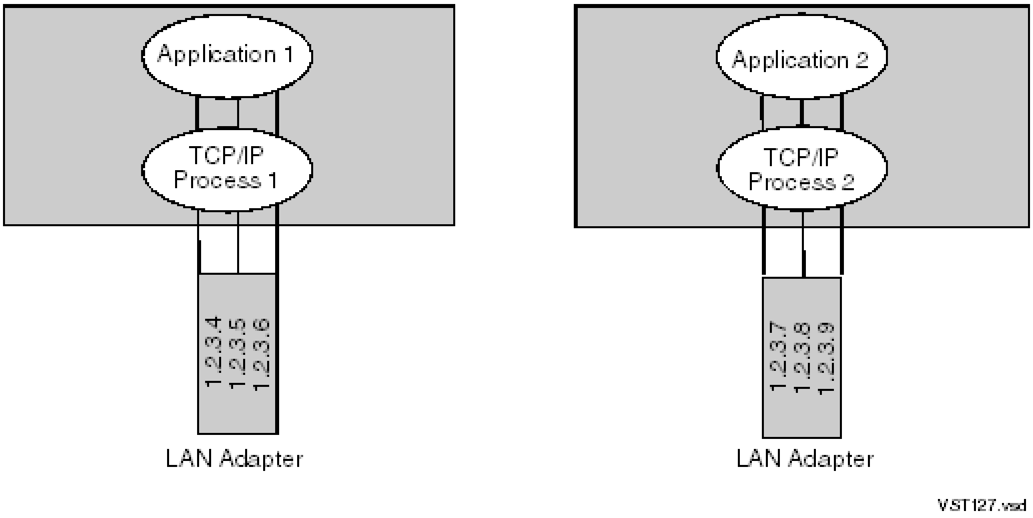

In NonStop TCP/IPv6, if you do not configure the environment to use logical-network partitioning and in CIP if you did not configure the environment to use Providers, applications using those subsystems cannot determine which interface they will get because the interface is no longer associated with the TCP/IP process used by the applications.

LNP support in NonStop TCP/IPv6 and Providers in CIP are similar to conventional TCP/IP in the sense that you can restrict an application's view of the network by associating the application with a TCP/IP process. LNP and Providers allow better control over an application's access to the network, by limiting the network resources available to the application to just those in the LNP or Provider that the application has been configured to use.

In NonStop TCP/IPv6, each LNP has its own set of IP addresses and SLSA logical interfaces (LIFs). In CIP, each Provider has its own set of IP addresses and Ethernet interfaces. An IP address used on one Provider or in one LNP cannot be used on a different Provider or different LNP, and an interface cannot be shared between Providers or LNPs. Applications on one Provider or LNP are isolated from applications on different Providers or LNPs on the same system in the same way they would be isolated if using different conventional TCP/IP processes. Communication between such applications is only possible through the attached local area networks. CIP and NonStop TCP/IPv6 do not forward packets between Providers or LNPs internally.

The difference between conventional TCP/IP and NonStop TCP/IPv6 with LNP or CIP with Providers is that the NonStop TCP/IPv6 transport process (TCP6SAM process) and the CIP transport process (CIPSAM), unlike the NonStop TCP/IP process, span all the processors in the whole system. The result is that an application in any processor, even when using LNP or Providers to restrict itself to specific interfaces, has direct access (with no interprocess, message-system hop) to the network adapter.

For more information about LNP, see the *TCP/IPv6 Configuration and Management Manual*. For more information about Providers, see the *Cluster I/O Protocols (CIP) Configuration and Management Manual*.

Most of the aspects of logical network partitioning that apply to NonStop TCP/IPv6 also apply to CIP with these differences:

- Network partitions are called Providers in CIP.
- In NonStop TCP/IPv6 you can assign individual interfaces to a partition whereas CIP, by default, assigns whole CLIMs, each of which is a group of five interfaces to a Provider. CIP allows individual interfaces to be assigned to a provider with the introduction of the MULTIPROV feature in J06.14 and H06.25.

Persistence

The NonStop TCP/IPv6 and CIP products on the Integrity NonStop systems participate in the system configuration database as generic processes and can be managed by the persistence manager. If you add their processes as generic processes, they start automatically upon system reload and restore their stored and subordinate objects.

For more information about configuring NonStop TCP/IPv6 to be persistent, see the *TCP/IPv6 Configuration and Management Manual*. For more information about configuring CIP to be persistent, see the *Cluster I/O Protcols (CIP) Configuration and Management Manual*. For more information about managing generic processes, see the *SCF Reference Manual for the Kernel Subsystem*.

What's Unique About Networking on Integrity NonStop BladeSystems

Networking on a NonStop BladeSystem offers the same features as networking on a NonStop NS-series system with these exceptions:

- The NonStop BladeSystem connects directly to the NonStop S-series I/O enclosure and supports the 6763 Common Communication ServerNet adapter (CCSA-2) with up to 4 SS7TE3 plug-in cards (PICs) and the token ring ServerNet adapter (TRSA). No other configurations of the CCSA are supported.
- Other differences in features exist. (See the *Cluster I/O Protcols (CIP) Configuration and Management Manual* for information about networking feature differences and see the *NonStop BladeSystem Planning Guide* for more information about features and differences of this system.)

NOTE: Only some Integrity NonStop systems support the S-series I/O enclosure. Check the planning guide for your system.

Summary Comparison of NonStop TCP/IP Products on H-Series and J-Series RVUs

This section compares at a high level, the three TCP/IP subsystems available on the H-series and J-series RVUs.

Table 1 Comparison of NonStop TCP/IP Products

	Conventional TCP/IP	NonStop TCP/IPv6	Cluster I/O Protocols (CIP)
Interface Types	• 10M / 100M / 1G Ethernet • ATM • SNAP (token-ring or Ethernet) • X.25	10M / 100M / 1G Ethernet	10M / 100M / 1G Ethernet
Jumbo Frames Supported	No	Yes	Yes

Table 1 Comparison of NonStop TCP/IP Products *(continued)*

	Conventional TCP/IP	NonStop TCP/IPv6	Cluster I/O Protocols (CIP)
Adapters	• MFIOB • IOAM+E4SA • IOAM+FESA • IOAM+GESA • IOAM+G4SA • VIO+G4SA • IOAM+ATM3SA • IOAM+TRSA • IOAM+CCSA	• MFIOB • IOAM+E4SA • IOAM+FESA • IOAM+GESA • IOAM+G4SA • VIO+G4SA	CLIM
Maximum Adapters	4	60	48
Maximum Hardware Interfaces	4 per process pair	240	240
IP versions	IPv4	IPv4, IPv6	IPv4, IPv6
SCTP Support	No	No	Yes
Protocol Stack Ancestry	BSD4.3	BSD/DEC	Linux
Protocol Stack Location	One NonStop Process Pair	One per NonStop system processor	Offloaded to CLIM
Remote Socket Support	Yes	No	No
Fault-Tolerant Sockets Support[1]	Yes	No	No
Round-Robin Listeners Support	No	Yes, maximum one per processor per port	Yes, no limit per processor per port
Network Partitioning Support	Yes	Yes	Yes
Maximum Partition Size	Interface	Interface	J06.14/H06.25 and later: Interface 06.13/H06.24 or earlier: CLIM
Minimum Partition Size	Interface	Interface	J06.14/H06.25 and later: Interface 06.13/H06.24 or earlier: CLIM
Interface Failover	None	Full	Partial
IPSec Support	No	No	Yes
Automatic DNS Updates	No	For IPv6	No

Table 1 Comparison of NonStop TCP/IP Products *(continued)*

	Conventional TCP/IP	NonStop TCP/IPv6	Cluster I/O Protocols (CIP)
Configuration Commands	• $*user-assigned name* through SCF • $ZZLAN through SCF	• $ZZTCP through SCF • $ZZLAN through SCF • TCP6SAM through SCF	• $ZZCIP through SCF • climconfig through CLIMCMD • CIPSAM through SCF • NonStop I/O Essentials plug-in to HP Systems Insight Manager (SIM)
Compatible Commands in SAM	N/A	• ABORT • INFO • LISTOPENS • NAMES • PRIMARY • STATS • STATUS • STOP • TRACE • VERSION	• ABORT • INFO • • NAMES • PRIMARY • • STATUS • STOP • TRACE • VERSION

[1] Fault-Tolerant Sockets transfer a socket from an application in one processor to its backup in another. Only Conventional TCP/IP supports them.

Unlike previous TCP/IP products, CIP does not use the SLSA subsystem. CIP uses a CLIM instead, which has higher speed and functionality than any of the SLSA-supported Ethernet adapters.

Table 2 Comparison of Communications Adapters and CLIMs

	E4SA	FESA	GESA	G4SA	VIO	CLIM
Number of Hardware Interfaces	4	1	1	4	8	5
Hardware Interface Speeds	10 Mbps	10/100 Mbps	10/100//1000 Mbps	• 2 10/100/1000 Mbps • 2 10/100 Mbps	• 2 10/100/1000 Mbps • 2 10/100 Mbps	5 10/100/1000 Mbps
Hardware Interface Media	4 copper	1 copper	1 copper	4 copper or 2 copper, 2 fiber	4 copper or 2 copper, 2 fiber	5 copper or 3 copper, 2 fiber, or 3 copper, 2 InfiniBand
Jumbo Frames	No	Yes	Yes	Yes	Yes	Yes
Onboard Processing	Filters	Filters	Filters	Filters	Filters	Full protocol stack

Previous TCP/IP products also support the multifunction I/O board (MFIOB), which is usually used for access to the dedicated service LAN. In addition to the five ports shown in Table 2, each CLIM has a separate port that is used for the dedicated service LAN.

2 Networking Concepts

This section provides a brief overview of networking concepts including the following topics:

- Address Resolution
- Name Resolution
- Allocation of IP Addresses (page 25)
- Routers (page 26)
- Switches (page 26)
- Gateways (page 26)
- Network Interface Name (page 26)
- Firewalls (page 27)
- "IP Security (IPSec)" (page 27)
- IPv6 (page 28)

For more detailed information about Internet concepts and services, see the *TCP/IP Configuration and Management Manual*.

Address Resolution

Address resolution refers to the mapping of IP addresses to lower-level addresses and is accomplished by the static binding of addresses or the dynamic binding of addresses. Static binding of addresses is used for the NonStop TCP/IP product over X.25. Dynamic binding of addresses is implemented with the Address Resolution Protocol (ARP) for IPv4 (defined in IEEE RFC 826) and is used for NonStop TCP/IP, NonStop TCP/IPv6 and CIP over Ethernet.

Address resolution with the dynamic binding of addresses for networks involves the use of the Neighbor Discovery Protocol. Using the Neighbor Discovery Protocol and stateless address auto configuration, an Integrity NonStop system configured as an IPv6 host discovers other nodes on the link, determines their link-layer addresses, finds routers, and maintains reachability information about the paths to active neighbors. See the *TCP/IPv6 Configuration and Management Manual* or *Cluster I/O Protocols (CIP) Configuration and Management Manual* for more information about the Neighbor Discovery Protocol and stateless address auto configuration in the IPv6 network. See also, IPv6 (page 28).

Name Resolution

For convenience, hosts are often referred to by name; in addition, for the world wide web, universal resource locators (URLs) locate a website location. The process of finding the IP address associated with either a host or a URL is a process of name-to-address mapping, is also called name resolution. There are two methods of resolving names, by:

- HOSTS File
- Domain Name System (DNS)

HOSTS File

A HOSTS file is an ASCII file on your system, by default in ZTCPIP (Guardian) or /etc (OSS) that lists the various host names associated with the IP addresses on the system. To use a HOSTS file for address resolution, you must configure the TCP/IP subsystem to use that HOSTS file; the default is for the TCP/IP subsystems is to use the Domain Name System (DNS) .

For procedures about configuring the various subsystems to use the HOSTS file, see the:

- *TCP/IPv6 Configuration and Management Manual*
- *TCP/IP Configuration and Management Manual*
- *Cluster I/O Protocols (CIP) Configuration and Management Manual*

Domain Name System (DNS)

The Internet has created an ever-increasing demand for IP addresses, and IP address management has presented a challenging task for administrators. In the past, administrators could manage the IP addresses in a single file containing all the host information (HOSTS File) with name-to-address mappings for every host connected to the network. Now assigning and maintaining new IP addresses and resolving domain names to IP addresses have become difficult and cumbersome tasks.

An effective solution to this problem is the Domain Name System (DNS), a distributed database that implements a name hierarchy for TCP/IP-based networks. DNS defines the rules for name syntax in a hierarchical name space and for delegation of authority over names. A name server is a server program that maps domain names to IP addresses. A set of DNS name servers operating at multiple sites cooperatively solve the domain name to IP address mapping problem.

Every time you use a domain name, a DNS service translates the name into the corresponding IP address. For example, the domain name www.sample.com translates to 188.135.212.3.

To use the domain name system, you must assign a name, in ARPANET standard format, to each system on the network or internetwork. You configure this name in your network configuration scripts. (See the *TCP/IPv6 Configuration and Management Manual* and the *Cluster I/O Protocols (CIP) Configuration and Management Manual* for details.)

You also need to create a set of ASCII files on each system which contains the addressing information the system needs. Instructions for creating these files are in the *TCP/IP Configuration and Management Manual*.

Once you have configured the domain name services, the network can access the node using the node's domain name and the domain name service routines will resolve the domain name to the node's IP address.

Allocation of IP Addresses

IP addresses are allocated in several ways. First, it is important to distinguish between IP addresses for components associated with a system and attached to the maintenance LAN (formerly known as the private LAN) and interfaces on components used for data communication between the system and other devices on the network. Devices in an Integrity NonStop system that are attached to the maintenance LAN include maintenance switches, NonStop System Consoles (NSCs), processor switches, ServerNet switch boards, CLIMs, and Uninterruptible Power Supply (UPS) units. Maintenance entity IP addresses come as statically configured IP addresses but can also be changed to dynamically configured IP addresses.

Address allocation has become more complex with the implementation of IPv6 (see IPv6 (page 28)) which introduced stateless address auto configuration, a process in which IP addresses actually expire, and with the increasing use of Dynamic Host Configuration Protocol (DHCP) servers.

IP address allocation for network interfaces on your NonStop systems can be obtained by contacting your Internet Service Provider (ISP) for an IPv6 address range for your site. See the IANA web page at:

`http://www.iana.org/ipaddress/ip-addresses.htm`

for more information about regional registries and address allocations.

Routers

A router is a device that has multiple network interfaces and transfers Internet Protocol (IP) packets from one network or subnet to another within an internetwork. (In many IP-related documents, this device is also referred to as a "gateway.")

Routing protocols find a path between network nodes. If multiple paths exist for a given protocol, the shorter paths are usually chosen. Each protocol has a cost or a metric that it applies to each path. In most cases, the lower the cost or metric for a given path, the more likely a protocol will choose it. In large local networks, there are often multiple paths to other parts of the local network. Routing daemons can be used to maintain near optimal routing to the other parts of the local network, and to recover from link failures in paths.

Switches

A switch receives messages from various devices on the network and routes the messages over the network to their appropriate destinations. The public telephone network provides the most obvious example of the use of switching, but switches are widely used in private networks as well.

Gateways

Networks that use different types of hardware and different protocols, such as TCP/IP and OSI, can communicate with each other through a gateway. Unlike a router, a gateway can translate the protocol of one network to a different protocol used by another network. When it is not being used to translate protocols, "gateway" is used interchangeably with "router".

Network Interface Name

A network interface is a communication device through which messages can be sent and received. Interface names on a G4SA are the logical interfaces (LIFs). On the CLIM, interface names are assigned by the operating system to the Ethernet interfaces. The names are eth1 through eth5. For diagrams showing the layout of the Ethernet interfaces on the CLIM, see the planning guide for your system. To determine the interface names when using NonStop TCP/IPv6 or NonStop TCP/IP, issue the SCF INFO LIF $ZZLAN.* command to the ServerNet LAN Systems Access (SLSA) subsystem and look at the "Name" column in the resulting display. For example:

```
1-> INFO LIF $ZZLAN.*
SLSA Info LIF
Associated
Name         Object    MAC Address         Type
$ZZLAN.CC10A CC1.0.A   00:00:00:00:00:00   WAN
$ZZLAN.LANY M0IE1.0.A  08:00:8E:00:7A:D9   Ethernet
$ZZLAN.LANX M0IE1.0.A  08:00:8E:00:7B:BA   Ethernet
$ZZLAN.L018 E0153.0.A  08:00:8E:00:78:3A   Ethernet
$ZZLAN.L019 E0153.0.B  08:00:8E:00:78:2D   Ethernet
$ZZLAN.L01A E0153.1.A  08:00:8E:00:78:2C   Ethernet
$ZZLAN.L01B E0153.1.B  08:00:8E:00:78:1C   Ethernet
$ZZLAN.L112I G1123.0.A 08:00:8E:00:97:B0   Ethernet
$ZZLAN.L01C T0154.0.A  08:00:8E:80:12:CE   Token Ring
$ZZLAN.L01A E0152.0.A  08:00:8E:AB:CD:EF   Ethernet
```

To determine the interface names when using CIP, issue the SCF STATUS CLIM DETAIL command. For example:

```
->STATUS CLIM $ZZCIP.* , DETAIL
.
.
.

Data Interface Status & IP Addresses:
 Flg Name      Status    LkP   Master / IP Family & Address
     eth5      UP        UP
                               IPv6: 3ffe:1200:190:1:21f:29ff:fe57:182e
                               IPv6: 3ffe:1200:190:2:21f:29ff:fe57:182e
                               IPv6: fe80::21f:29ff:fe57:182e
     eth4      UP        UP
                               IPv6: 3ffe:1200:190:1:21f:29ff:fe57:182f
                               IPv6: 3ffe:1200:190:2:21f2:9ff:fe57:182f
                               IPv6: fe80::21f:29ff:fe57:182f
     eth3      UP        UP
                               IPv6: 3ffe:1200:190:1:21f:29ff:fe57:182c
                               IPv6: 3ffe:1200:190:2:21f:29ff:fe57:182c
                               IPv6: fe80::21f:29ff:fe57:182c
     eth2      UP        UP
                               IPv6: 3ffe:1200:190:2:21f:29ff:fe57:182d
                               IPv6: 3ffe:1200:190:1:21f:29ff:fe57:182d
                               IPv6: fe80::21f:29ff:fe57:182d
.
.
.
```

Firewalls

A firewall is a system or group of systems that enforces an access control policy between two or more networks. The actual means by which access control is accomplished varies widely, but in principle, the firewall can be thought of as a pair of mechanisms: One which exists to block traffic, and the other which exists to permit traffic. Some firewalls place a greater emphasis on blocking traffic, while others emphasize permitting traffic. The most important thing to recognize about a firewall is that it implements an access control policy. Before you install a firewall, you need to know what kind of access you want to allow or deny. Also note that because the firewall is a mechanism for enforcing policy, it imposes its policy on everything behind it. Administrators for firewalls managing the connectivity for a large number of hosts have a heavy responsibility.

IP Security (IPSec)

Although IPSec is only supported in the CIP product, it is not supported on the IB CLIM.

The IP security architecture (IPSec) defines basic security mechanisms at the network level so they can be available to all the layered applications. The security techniques adopted in IPSec have been designed to be easily inserted in both IPv4 and IPv6.

IPSec security services are offered by means of two dedicated extension headers, the Authentication Header (AH) and the Encapsulating Security Payload (ESP), and through the use of cryptographic key management procedures and protocols.

The AH header was designed to ensure authenticity and integrity of the IP packet. It also provides an optional anti-replay service. Its presence guards against illegal modification of the IP fixed fields, packet spoofing and, optionally, against replayed packets. On the other hand, the ESP header provides data encapsulation with encryption to ensure that only the destination node can read the payload conveyed by the IP packet. ESP may also provide packet integrity and authenticity, and an anti-replay service. The two headers can be used separately or they can be combined to provide the desired security features for IP traffic.

Each header can be used in one of the two defined modalities: transport mode and tunnel mode. While in transport mode the security headers provide protection primarily for upper layer protocols, in tunnel mode the headers are applied to tunneled IP packets, thus providing protection to all fields of the original IP header.

Both AH and ESP exploit the concept of security association (SA) to agree upon the security algorithms, transforms and parameters shared by the sender and the receiver of a protected traffic flow. Each IP node manages a set of SAs, with at least one SA for each secure communication. The SAs currently active are stored inside a database, known as the security association database (SAD). An entry in the SAD (for example, a security association) is uniquely identified by a triplet consisting of a security parameter index (SPI), an IP destination address, and a security protocol (AH or ESP) identifier. The security parameter index (SPI) is transmitted inside both the AH and

ESP headers since it is used to choose the right SA to be applied for decrypting and authenticating the packet.

In unicast transmissions, the SPI is normally chosen by the destination node and sent back to the sender when the communication is set up. In multicast transmissions, the SPI must be common to all the members of the multicast group. Each node must be able to correctly identify the right SA by combining the SPI with the multicast address. The negotiation of a SA (and the related SPI) is an integral part of the protocol for the exchange of security keys.

Specific security requirements are defined at each node usually by means of an ordered list of admission rules (or policies), which form the node's security policy database (SPD). The protection provided to each incoming and outgoing traffic flow is verified by consulting the SPD. In general, packets are selected for one of three processing modes based on IP and transport layer header information matched against entries in the SPD. Each packet is either afforded IPSec security services, discarded, or allowed to bypass IPSec, based on the applicable policies found in the database.

IPv6

The IPv6 protocol extends the IP address to 128 bits compared to the 32 bits of IPv4 addresses. The NonStop TCP/IPv6 subsystem provides IPv6 functionality on the Integrity NonStop system using three modes of operation: pure IPv6, in which the system supports only IPv6 communications, DUAL, in which the system supports both IPv4 and IPv6 communications, and pure IPv4, in which the system provides only IPv4 communications. The CIP subsystem also supports IPv6 with an attribute in the Provider object, which can be set to either INET (IPv4) or DUAL.

Much of the Internet consists of IPv4 networks; an IPv6-enabled system can communicate across IPv4 networks by using tunneling. IPv6 tunneling requires an IPv4 address on both ends of the communication; IPv6 packets are then encapsulated in IPv4 packets so that they can be transmitted across an IPv4 network. The IPv6-aware host or router decapsulates the IPv6 datagrams, forwarding them as needed. IPv6 tunneling eases IPv6 deployment by maintaining compatibility with the large existing base of IPv4 hosts and routers. Figure 9 depicts an IPv6 tunneling scenario.

Figure 9 IPv6 Tunneling

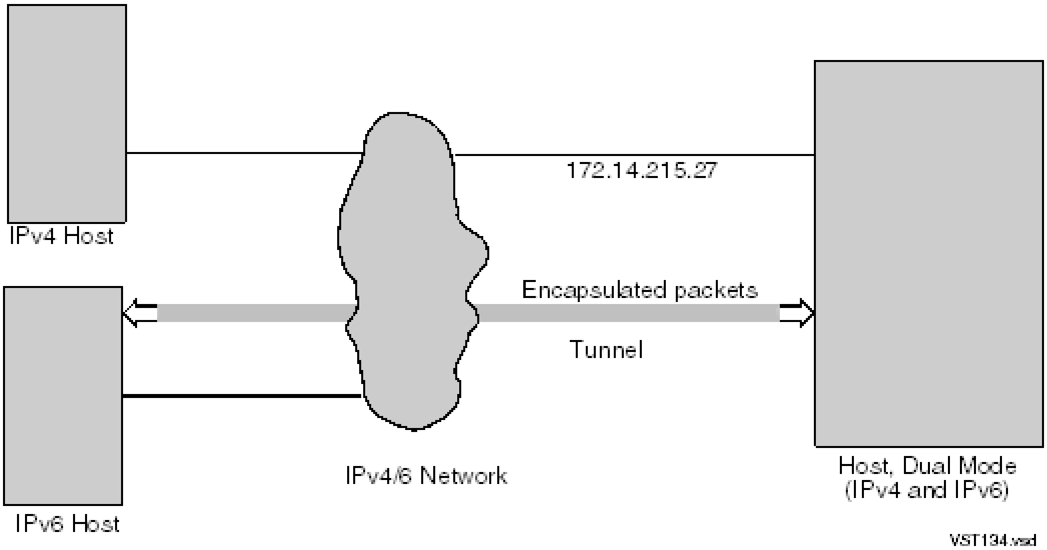

For more information about IPv6, see the *TCP/IPv6 Configuration and Management Manual*.

3 Planning for Migrating Networking Solutions to H-Series and J-Series RVUs

This section provides examples of product-suite requirements for some legacy networking solutions on H-series and J-series RVUs. These examples were chosen to show the general placement and relationships of legacy networking software and hardware in the H-series and J-series RVU environments.

In addition, this section provides guidance for finding Internet application documentation.

Topics covered in this section include:

- Planning Your Networking Solutions
- Migrating From a Platform Other Than a NonStop Server (page 33)
- Internet Application Product Documentation (page 33)

Planning Your Networking Solutions

This subsection provides information about which subsystems, products and manuals you might need for various networking solutions on the Integrity NonStop system, including:

- SNAX/XF (Includes Function From SNAX/APN)
- X.25 Communications (page 31)
- Asynchronous Wide Area Network (AWAN) Connectivity (page 32)

In addition, this subsection addresses Selecting Your TCP/IP Product (page 33)

SNAX/XF (Includes Function From SNAX/APN)

To use SNAX/XF on H-series RVUs and J-series RVUs, you need these hardware products and subsystems:

- SNAX/XF
- NonStop TCP/IP, NonStop TCP/IPv6 or CIP
- Port Access Method (PAM) or Wide Area Network WAN subsystem
- SLSA (unless using CIP)
- CLIM, G4SA, FESA, GESA, E4SA, or TRSA
- NonStop S-series I/O enclosure (optional, used for TRSA, E4SA, FESA, CCSA, or GESA connectivity)

Figure 10 Products for SNAX/XF

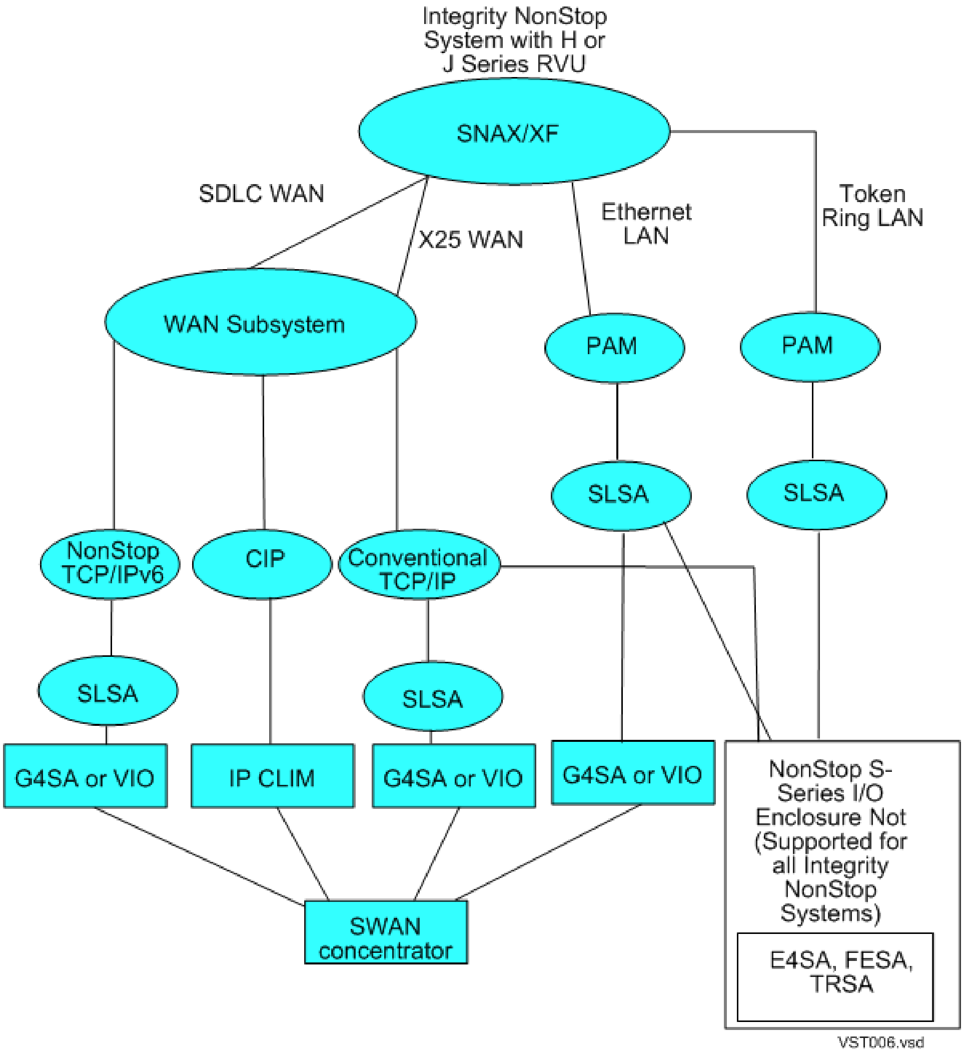

The manuals you might need for running SNAX/XF include:

- *Cluster I/O Protocols (CIP) Configuration and Management Manual*
- *Ethernet Adapter Installation and Support Guide*
- *Gigabit Ethernet 4-Port Adapter Installation and Support Guide*
- *Gigabit Ethernet Adapter Installation and Support Guide*
- *Introduction to Networking for HP NonStop S-Series Servers*
- *LAN Configuration and Management Manual*
- Planning guide for your system
- *Port Access Method (PAM) Configuration and Management Manual*
- *SNAX/XF and SNAX/APN Configuration and Management Manual*
- *Token Ring Adapter Installation and Support Guide*
- *WAN Subsystem Configuration and Management Manual*

X.25 Communications

For X.25 communications on H-series or J-series RVUs, you need:

- X25AM
- CIP, NonStop TCP/IP, or NonStop TCP/IPv6
- WAN subsystem
- SLSA (for NonStop TCP/IP and NonStop TCP/IPv6 only)
- CLIM, G4SA, E4SA, FESA, or GESA
- A SWAN or SWAN 2 Concentrator
- NonStop S-series I/O enclosure (optional, used for TRSA, E4SA, FESA, or GESA connectivity)

Figure 11 Products for X.25 Communications for H-Series or J-Series RVUs

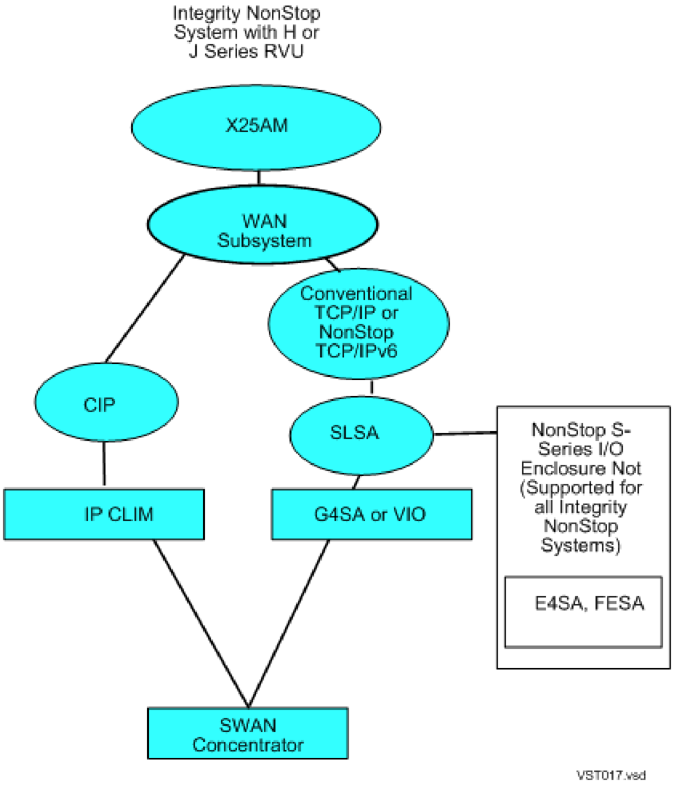
Integrity NonStop System with H or J Series RVU
X25AM
WAN Subsystem
Conventional TCP/IP or NonStop TCP/IPv6
CIP
SLSA
NonStop S-Series I/O Enclosure Not (Supported for all Integrity NonStop Systems)
IP CLIM
G4SA or VIO
E4SA, FESA
SWAN Concentrator
VST017.vsd

The manuals you might need for X.25 communications may include:

- *Cluster I/O Protocols (CIP) Configuration and Management Manual*
- *Ethernet Adapter Installation and Support Guide*
- *Gigabit Ethernet 4-Port Adapter Installation and Support Guide*
- *Gigabit Ethernet Adapter Installation and Support Guide*
- *Introduction to Networking for HP NonStop S-Series Servers*
- *LAN Configuration and Management Manual*
- Planning guide for your system
- *Port Access Method (PAM) Configuration and Management Manual*
- *SWAN Concentrator Installation and Support Guide*
- *SWAN 2 Concentrator Installation and Support Guide*

- *WAN Subsystem Configuration and Management Manual*
- *X25AM Configuration and Management Manual*

Asynchronous Wide Area Network (AWAN) Connectivity

If you want AWAN connectivity on the H-series or J-series RVU you need:

- Telserv
- FASTPTCP print processes
- CIP, NonStop TCP/IP, NonStop TCP/IPv6
- SLSA (for NonStop TCP/IP or NonStop TCP/IPv6 only)
- CLIM or G4SA
- AWAN 3886 server
- NonStop S-series I/O enclosure (optional, used for TRSA, E4SA, FESA, or GESA connectivity)

NOTE: AWAN is not supported on NonStop NS2000, NS2100, and NS2200-series systems.

Figure 12 Products for AWAN Connectivity on H-Series and J-Series RVUs

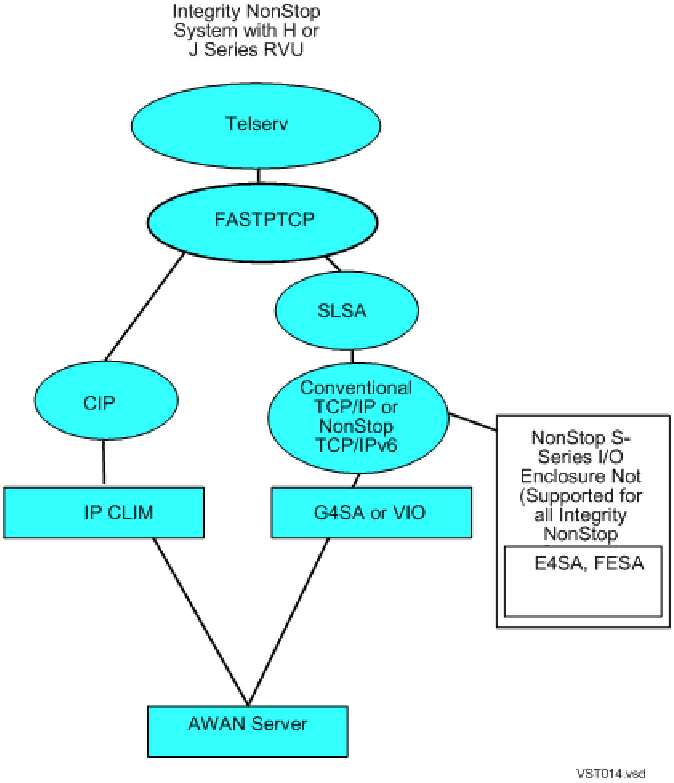
Integrity NonStop System with H or J Series RVU
Telserv
FASTPTCP
SLSA
CIP
Conventional TCP/IP or NonStop TCP/IPv6
NonStop S-Series I/O Enclosure Not (Supported for all Integrity NonStop
E4SA, FESA
IP CLIM
G4SA or VIO
AWAN Server
VST014.vsd

Manuals that you might need for AWAN connectivity include:

- *AWAN 3886 Server Installation and Configuration Guide*
- *Cluster I/O Protocols (CIP) Configuration and Management Manual*
- *Ethernet Adapter Installation and Support Guide*
- *Fast Ethernet Adapter Installation and Support Guide*
- *Gigabit Ethernet 4-Port Adapter Installation and Support Guide*
- *Gigabit Ethernet Adapter Installation and Support Guide*
- *Introduction to Networking for HP NonStop S-Series Servers*

- *LAN Configuration and Management Manual*
- *LAN Configuration and Management Manual*
- *Spooler FASTP Network Print Processes Manual*
- *TCP/IP Configuration and Management Manual*
- *TCP/IPv6 Configuration and Management Manual*
- *Telserv Manual*

Selecting Your TCP/IP Product

HP recommends using NonStop TCP/IPv6 or CIP for your TCP/IP needs because of the superior processing power of the parallel architecture that underlies these products. (See Integrity NonStop System Networking Compared to NonStop S-Series Server Networking (page 14), the *TCP/IPv6 Configuration and Management Manual*, and the *Cluster I/O Protocols (CIP) Configuration and Management Manual*.) Because you can use NonStop TCP/IPv6 and CIP for IPv6 communications (see IPv6 (page 28)) you can use these products to gain the benefits of the architecture without implementing IPv6 communications in your environment. If you want to use IPv6 communications, you must use NonStop TCP/IPv6 or CIP.

To use any data communication protocols other than Ethernet or X.25, you must use NonStop TCP/IP instead of NonStop TCP/IPv6 or CIP. NonStop TCP/IP provides access to these:

- Token Ring
- ATM

Migrating From a Platform Other Than a NonStop Server

If you are migrating to the Integrity NonStop system from another platform and do not have NonStop S-series systems in your environment, your networking options will be restricted to Gigabit Ethernet and you should plan to migrate your applications to run over TCP/IP, ServerNet WAN, or CLIM connections.

If you are migrating to the Integrity Nonstop NS-series system from the NonStop S-series system, you can retain most of your existing networking solutions, including token ring and ATM, but you will not be able to use IPX/SPX or ServerNet/FX directly. (However, ServerNet/FX can provide inter connectivity between NonStop S-series systems and NonStop K-series systems in an environment that includes Integrity NonStop systems. See Figure 3 (page 12).) If you are going to use token ring or ATM, you can run your applications on the Integrity NonStop NS-series system and manage the subsystems for each of those protocols on the Integrity NonStop NS-series system as well. You do not need to do anything to your applications to make them run on the Integrity NonStop NS-series system but your hardware must include a connection from the IOAM to the NonStop S-series I/O enclosure.

The option of attaching to the NonStop S-series I/O enclosure is only available for NonStop BladeSystems for the 6763 Common Communication ServerNet Adapter (CCSA-2). This support was added to provide the SS7 functionality only available through the CCSA.

If you are using NonStop TCP/IPv6 and want to tune it by configuring the QIO subsystem to run in the global privileged space, you can take advantage of the Integrity NonStop system's increase of this memory location to 256 MB. See the *QIO Configuration and Management Manual* for information about how and when to make this adjustment.

Internet Application Product Documentation

You can find general information about the products listed in Table 3 in the *Introduction to Networking for HP NonStop S-Series Servers.*

Table 3 Internet Application Product Documentation

Product	Manuals
BEA WebLogic Server for the NonStop Server	*HP NonStop Server Platform Guide for WebLogic Server 8.1*
iTP Secure WebServer	*iTP Secure WebServer System Administrator's Guide*
iTP Active Transaction Pages (ATP)	*iTP Active Transaction Pages (iTP ATP) Programmer's Guide*
NonStop Servlets for JavaServer Pages (NSJSP)	*NonStop Servlets for JSP System Administrator's Guide*
NonStop SOAP	*NonStop SOAP User's Manual*
NonStop SOAP for Java	*NonStop SOAP for Java User's Manual*

Index