Product Bulletin - Public #487

Cisco IOS Software Release 11.2

Introduction

This Product Bulletin describes new features introduced in Cisco IOS™ software release 11.2. Please refer to Product Bulletin #488 for ordering procedures, as well as a summary of platforms supported and platform-specific features. Product Bulletin #488 will be available upon First Commercial Shipment (FCS) of Cisco IOS software release 11.2.

1 Routing Protocols

1.1 IP Protocol and Feature Enhancements

1.1.1 On Demand Routing

Description: On Demand Routing (ODR) is a mechanism that provides minimum-overhead IP routing for stub sites. The overhead of a general dynamic routing protocol is avoided, without incurring the configuration and management overhead of using static routing.

A stub router is the peripheral router in a hub and spoke network topology. Stub routers commonly have a WAN connection to the hub router and a small number of LAN network segments (stub networks) that are connected directly to the stub router. To provide full connectivity, the hub routers can be statically configured to know that a particular stub network is reachable via a specified access router. However, if there are multiple hub routers, many stub networks, or asynchronous connections between hubs and spokes, the overhead required to statically configure knowledge of the stub networks on the hub routers becomes too great.

ODR simplifies installation of IP stub networks in which the hub routers dynamically maintain routes to the stub networks. This is accomplished without requiring the configuration of an IP routing protocol at the stub routers. With ODR, the stub advertises IP prefixes corresponding to the IP networks that are configured on its directly connected interfaces. Because ODR advertises IP prefixes, rather than IP network numbers, ODR is able to carry Variable Length Subnet Mask (VLSM) information.

Once ODR is enabled on a hub router, the router begins installing stub network routes in the IP forwarding table. The hub router can also be configured to redistribute these routes into any configured dynamic IP routing protocols. IP does not need to be configured on the stub router. With ODR, a router is automatically considered to be a stub when no IP routing protocols have been configured on it.

Benefits: ODR minimizes the configuration and bandwidth overhead required to provide full routing connectivity. Moreover, it eliminates the need to configure an IP routing protocol at the stub routers.

Considerations: The routing protocol which ODR generates is propagated between routers using Cisco's Discovery Protocol (CDP). Thus, ODR is partially controlled by the configuration of CDP. Specifically,

- If CDP is disabled, the propagation of ODR routing information will cease.
- By default, CDP sends updates every 60 seconds. This update interval may not be frequent enough to provide fast reconvergence of IP routers on the hub router side of the network. A faster reconvergence rate may be necessary if the stub connects to several hub routers via asynchronous interfaces (such as modem lines).

- ODR may not work well with dial-on demand routing (DDR) interfaces, as CDP packets will not cause a DDR connection to be made.

It is recommended that IP filtering be used to limit the network prefixes which the hub router will permit to be learned dynamically through ODR. If the interface has multiple logical IP networks configured (via the IP secondary command), only the primary IP network is advertised through ODR.

Product Marketing Contact: Martin McNealis

1.1.2 Open Shortest Path First (OSPF) Enhancements

1.1.2.1 OSPF On-Demand Circuit

Description: OSPF On-Demand Circuit is an enhancement to the OSPF protocol, as described in RFC 1793, that allows efficient operation over demand circuits such as ISDN, X.25 SVCs, and dial-up lines. Previously, the period nature of OSPF routing traffic mandated that the underlying data-link connection needed to be open constantly, resulting in unwanted usage charges. With this feature, OSPF Hellos and the refresh of OSPF routing information is suppressed for on-demand circuits (and reachability is presumed), allowing the underlying data-link connections to be closed when not carrying application traffic.

Figure 1. OSPF On-Demand Circuit

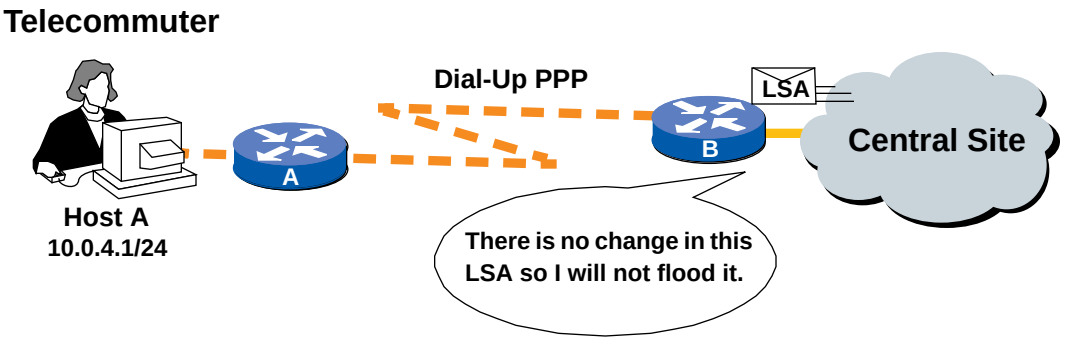

Benefits: OSPF On-Demand Circuit is ideally suited to network infrastructures with an OSPF backbone and peripheral sites connecting to the central backbone network. The feature allows the consolidation on a single routing protocol and the benefits of the OSPF routing protocol across the entire network, without incurring excess connection costs.

OSPF On-Demand Circuit is also applicable when bandwidth is at a premium, for example in networks with low-speed links.

Considerations: If the router is part of a point-to-point topology, only one end of the demand circuit needs to be configured for OSPF On-Demand Circuit operation. In point-to-multipoint topologies, all appropriate routers must be configured with OSPF On-Demand Circuit. All routers in an area must support this feature -- that is, be running Cisco IOS software release 11.2 or greater.

Product Marketing Contact: Martin McNealis

1.1.2.2 OSPF Not-So-Stubby Areas (NSSA)

Description: As part of the OSPF protocol's support for scalable, hierarchical routing, peripheral portions of the network can be defined as "stub" areas, so that they do not receive and process external OSPF advertisements. Stub areas are generally defined for low end routers with limited memory and CPU, which have low-speed connections, and are in a default route configuration.

OSPF Not-So-Stubby-Areas (NSSA) defines a more flexible, hybrid method, whereby stub areas can import external OSPF routes in a limited fashion, so that OSPF can be extended across the stub to backbone connection

Figure 2. OSPF Not-So-Stubby Areas.

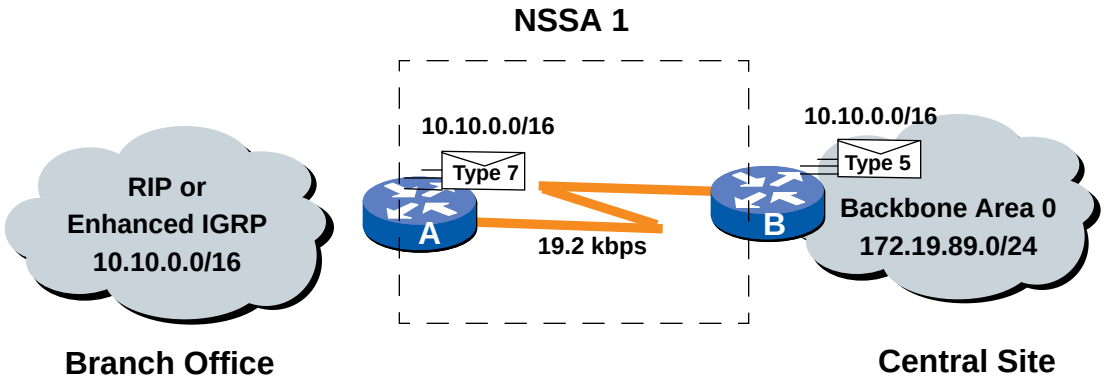

Benefits: NSSA enables OSPF to be extended across a stub area to backbone area connection to become logically part of the same network. As with regular stub areas, this can be deployed in conjunction with multiple Area Border Routers (ABRs) and NSSA ABRs can inject a default route into the NSSA.

Considerations: Since OSPF's advertisements indicating topology changes are flooded over On-Demand Circuits, it is recommended that demand circuits within OSPF stub areas or within NSSAs are isolated from topology changes.

Product Marketing Contact: Martin McNealis

1.1.3 Border Gateway Protocol version 4 (BGP4) Enhancements

1.1.3.1 BGP4 Soft Configuration

Description: BGP4 soft configuration allows BGP4 policies to be configured and activated without clearing the BGP session, hence without invalidating the forwarding cache. This enables policy reconfiguration without causing short-term interruptions to traffic being forwarded in the network.

Benefits: This prevents periods of instability within an Autonomous System (AS) when administrators need to modify their policies.

Considerations: Outbound reconfiguration is preferable over inbound, since changes in inbound policy are memory intensive -- the router must remember all updates received from a neighbor.

Product Marketing Contact: Martin McNealis

1.1.3.2 BGP4 Multipath Support

Description: BGP4 Multipath support provides BGP load balancing between multiple Exterior BGP (EBGP) sessions. If there are multiple EBGP sessions between the local Autonomous System (AS) and the neighboring AS, multipath support allows BGP to load balance among these sessions. Depending on the switching mode, per packet or per destination load balancing is performed.

Benefits: BGP4 Multipath support provides greater overall redundancy.

Considerations: BGP4 Multipath Support can support up to six paths.

Product Marketing Contact: Martin McNealis

1.1.3.3 BGP4 Prefix Filtering with Inbound Route Maps

Description: This allows prefix-based matching support to the inbound neighbor route map. With this addition, an inbound route map can be used to enforce prefix-based policies.

Benefits: Prefix Filtering with Inbound Route Maps allows the network administrator to specify the level of summarization or aggregation that will be accepted in an advertised network prefix from a neighboring Autonomous System.

Considerations: none.

Product Marketing Contact: Martin McNealis

1.1.4 Network Address Translation

Description: Network Address Translation (NAT) provides a mechanism for a privately addressed network to access registered networks, such as the Internet, without requiring a registered subnet address. This eliminates the need for host renumbering and allows the same IP address range to be used in multiple intranets.

With NAT, the privately addressed network (designated as “inside”) continues to use its existing private or obsolete addresses. These addresses are converted into legal addresses before packets are forwarded onto the registered network (designated as “outside”). The translation function is compatible with standard routing; the feature is required only on the router connecting the inside network to the outside domain.

Translations can be static or dynamic in nature. A static address translation establishes a one-to-one mapping between the inside network and the outside domain. Dynamic address translations are defined by describing the local addresses to be translated and the pool of addresses from which to allocate outside addresses. Allocation is done in numeric order and multiple pools of contiguous address blocks can be defined.

Figure 3. Network Address Translation.

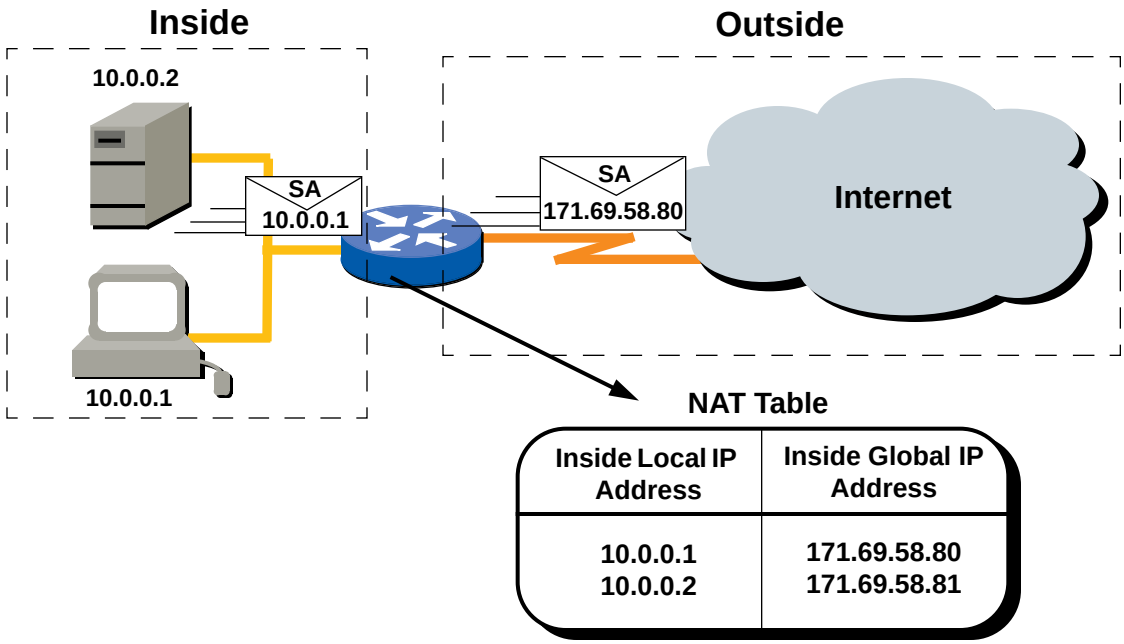

Benefits: NAT provides the following benefits:

- Eliminates readdressing overhead. NAT eliminates the need to readdress all hosts that require external access, saving time and money.
- Conserves addresses through application port-level multiplexing. With NAT, internal hosts can share a single registered IP address for all external communications. In this type of configuration, relatively few external addresses are required to support many internal hosts, thus conserving IP addresses.
- Protects network security. Because private networks do not advertise their addresses or internal topology, they remain reasonably secure when used in conjunction with NAT to gain controlled external access.
- Since the addressing scheme on the inside network may conflict with registered addresses already assigned within the Internet, NAT can support a separate address pool for overlapping networks and translate as appropriate.

Considerations: Applications that use raw IP addresses as a part of their protocol exchanges are incompatible with Network Address Translation. Typically, these are less common applications that do not use fully qualified domain names.

Product Marketing Contact: Martin McNealis

1.1.5 Named IP Access Control List

Description: The Named IP Access Control List (ACL) feature gives network managers the option of using names for their access control lists. Named IP Access Control Lists function similarly to their numbered counter-parts, except that they use names instead of numbers.

This feature also includes a new configuration mode, which supports addition and deletion of single lines in a multi-line access control list.

Benefits: This feature eliminates some of the confusion associated with maintaining long access control lists. Meaningful names can be assigned, making it easier to remember which service is controlled by which access control list. Moreover, this removes the limit of 100 extended and 99 standard access control lists, so that additional IP access control lists can be configured.

The new configuration feature allows a network manager to edit access control lists, rather than recreating the entire list. This saves time and eliminates the potential risk of subjecting the network to compromise.

Considerations: Currently, only packet and route filters can use Named IP Access Control Lists. Also, named IP Access Control Lists are not backward-compatible with earlier releases of Cisco IOS software.

Named IP Access Control Lists are not currently supported with Distributed Fast Switching.

Product Marketing Contact: Peter Long

1.2 Multimedia and Quality of Service

1.2.1 Resource Reservation Protocol

Description: Resource Reservation Protocol (RSVP) enables applications to dynamically reserve necessary network resources from end-to-end for different classes of service. An application, which acts as a receiver for a traffic stream, initiates a request for reservation of resources (bandwidth) from the network, based on the application's required quality of service. The first RSVP-enabled router that receives the request informs the requesting host whether the requested resources are available or not. The request is forwarded to the next router, towards the sender of the traffic stream. If the reservations are successful, an end-to-end pipeline of resources is available for the application to obtain the required quality of service. RSVP enables applications with real-time traffic needs, such as multimedia applications, to coexist with bursty applications on the same network. RSVP works with both unicast and multicast applications.

Figure 4. Resource Reservation Protoco

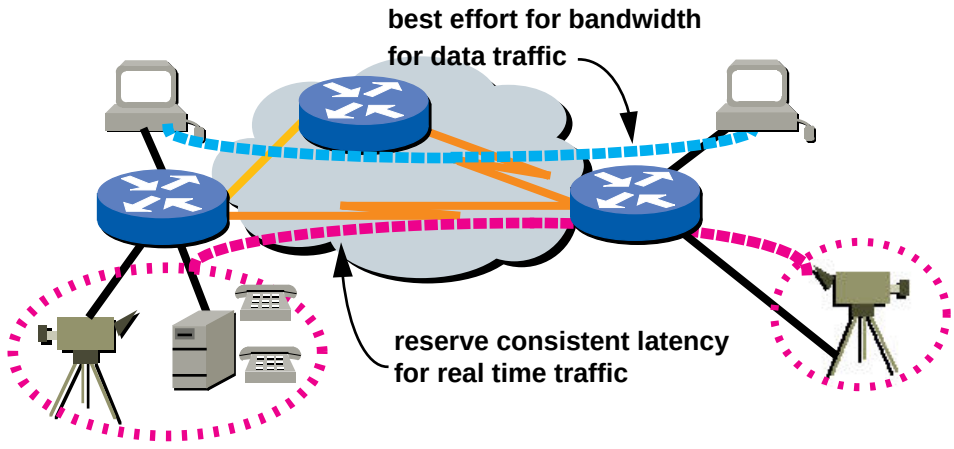

Benefits: Multimedia applications can coexist with bursty applications on the same network. Parallel networks are not required to support different applications; new multimedia applications can be supported while continuing to support legacy applications.

RSVP protects customers' investments. Existing networks can be upgraded through software to provide guaranteed quality of service for multimedia applications, eliminating the need for an expensive "fork-lift" upgrade (hardware replacement) to support multimedia applications.

Considerations: RSVP requires both a network implementation and a client implementation. Applications need to be RSVP-enabled to take advantage of RSVP functionality. Currently, Precept provides an implementation of RSVP for Windows-based PCs. Companies such as Sun and Silicon Graphics have demonstrated RSVP on their platforms. Several application developers are planning to take advantage of RSVP in their applications.

Product Marketing Contact: Seenu Banda

1.2.2 Random Early Detection

Description: Random Early Detection (RED) helps eliminate network congestion during peak traffic loads. RED uses the characteristics of robust transport protocols (such as TCP) to reduce transmission volume at the source when traffic volume threatens to overload a router's buffer resources.

The effectiveness of RED depends on the protocol mix in the network. In networks with a high percentage of TCP (or other robust protocol) traffic, RED is an effective way to alleviate congestion. RED "throttles back" lower priority traffic first, allowing higher priority traffic (as designated by an RSVP reservation or the IP precedence value) to continue unabated.

Benefits: RED works with RSVP to maintain end-to-end quality of service during peak traffic loads. Congestion is avoided by selectively dropping traffic during peak load periods. This is performed in a manner designed to damp out waves of sessions going through TCP slow start.

Existing networks can be upgraded to better handle RSVP and priority traffic. Additionally, RED can be used in existing networks to manage congestion more effectively on higher speed links where fair queuing is expensive.

Considerations: RED is not recommended for networks with a significant percentage of protocols that are not robust in the face of packet loss (such as AppleTalk or Novell NetWare.)

RED is a queuing technique; it cannot be used on the same interface as other queuing techniques, such as Standard Queuing, Custom Queuing, Priority Queuing, or Fair Queuing.

Product Marketing Contact: Erik Gilbert

1.2.3 Generic Traffic Shaping

Description: Generic Traffic Shaping (also called Interface Independent Traffic Shaping) helps reduce the flow of outbound traffic from a router interface into a backbone transport network when congestion is detected in the downstream portions of the backbone transport network or in a downstream router. Unlike the Traffic Shaping over Frame Relay features (see 3.2.2), which are specifically designed to work on interfaces to Frame Relay networks, Generic Traffic Shaping works on interfaces to a variety of layer 2 data link technologies (including Frame Relay, SMDS, Ethernet, etc.)

Figure 5. Generic Traffic Shaping

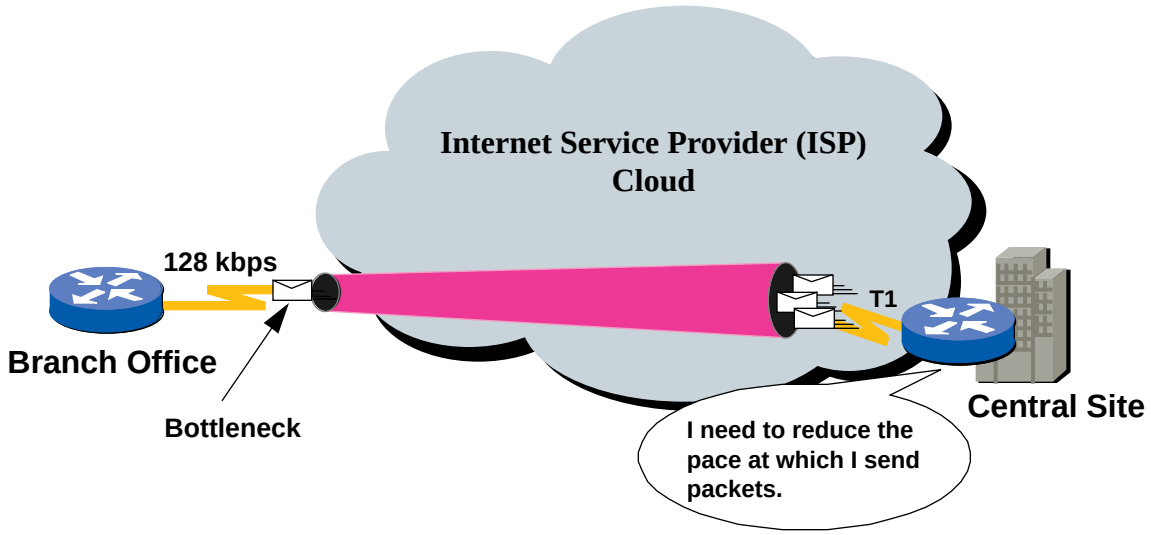

Benefits: Topologies that have high speed links (e.g., a central site) feeding into lower speed links (e.g., at remote or branch sites) often experience bottlenecks at the remote end because of the speed mismatch. Generic Traffic Shaping helps eliminate the bottleneck situation by throttling back traffic volume at the source end.

Routers can be configured to transmit at a lower bit rate than the interface bit rate. Service providers or large enterprises can use the feature to partition, for example, T1 or T3 links into smaller channels to match service ordered by customers.

Packet loss in the service provider's network can be limited by throttling the traffic back at the source, thus improving service predictability.

Generic Traffic Shaping protects a customer's investment. Existing routed protocol (e.g., IP, IPX, SNA) networks can be upgraded to provide consistent quality of service for RSVP and priority traffic across backbone transport networks (e.g., carrier networks). Existing routed protocol networks can be upgraded to act in concert with transport networks to provide better end-to-end service. Interface density on existing networks can be increased by limiting traffic on each interface to a predetermined desired level.

Considerations: Generic Traffic Shaping implements a Weighted Fair Queuing (WFQ) on an interface or sub-interface to allow the desired level of traffic flow. The feature consumes router memory and CPU resources, and so must be used judiciously to regulate critical traffic flows while not degrading overall router performance.

Product Marketing Contact: Erik Gilbert

1.3 Multiprotocol Routing

1.3.1 Enhanced IGRP Optimizations

Description: With the wide-scale deployment of Enhanced Interior Gateway Routing Protocol (Enhanced IGRP) in increasingly large and complex customer networks, Cisco has been able to continuously monitor and refine Enhanced IGRP operation, integrating several key optimizations. Optimizations have been made in the allocation of bandwidth, use of processor and memory resources, and mechanisms for maintaining information about peer routers.

Benefits:

- **Intelligent Bandwidth Control --** In network congestion scenarios, packet loss, especially the dropping of routing protocol messages, adversely affects convergence time and overall stability. To prevent this problem, Enhanced IGRP now takes into consideration the available bandwidth (at a granularity of per subinterface/virtual circuit if appropriate) when determining the rate at which it will transmit updates. Interfaces can also be configured to use a certain (maximum) percentage of the bandwidth, so that even during routing topology computations, a defined portion of the link capacity remains available for data traffic.
- **Improved Processor and Memory Utilization --** Enhanced IGRP derives the distributed routing tables from topology databases that are exchanged between peer routers. This CPU computation has now been made significantly more efficient as has the protocol's queuing algorithm, resulting in improved memory utilization. The combination of these factors further increases Enhanced IGRP's suitability for deployment, particularly on low-end routers.
- **Implicit Protocol Acknowledgments** -- Enhanced IGRP running within a router maintains state and reachability information about other neighboring routers. This mechanism has been modified so that it no longer requires explicit notifications to be exchanged but rather will accept any traffic originating from a peer as a valid indication that the router is operational. This provides greater resilience under extreme load.
- **IPX Service Advertisement Interleaving --** Large IPX environments are typically characterized by many Service Advertisements which can saturate lower speed links at the expense of routing protocol messages. Enhanced IGRP now employs an interleaving technique to ensure that both traffic types receive sufficient bandwidth in large IPX networks.

Considerations: These key enhancements are particularly applicable in networking environments having many low-speed links (typically in "hub and spoke" topologies), in Non-Broadcast-Multiple-Access (NBMA) wide-area networks such as Frame Relay, ATM, or X.25 backbones and in highly redundant, dense router-router peering configurations. It should be noted that the basic Enhanced IGRP routing algorithm that exhibits very fast convergence and guaranteed loop-free paths has not changed, so there are no backwards compatibility issues with earlier versions.

Product Marketing Contact: Martin McNealis

1.4 Switching Features

1.4.1 Integrated Routing and Bridging

Description: Cisco delivers another integral component of the Cisco Fusion architecture for scaleable and efficient integration of layer 2 and layer 3 environments in multiprotocol networks. Integrated Routing and Bridging (IRB) delivers the functionality to extend vLANs and layer 2 bridged domains across the groups of interfaces on Cisco IOS software-based routers and to interconnect them to the routed domains within the same router.

Concurrent Routing and Bridging (CRB) supports routing and bridging of the same layer 3 protocol on multiple independent sets of interfaces within the same Cisco IOS software-base router. IRB extends CRB by providing the capability to route between bridged and routed interfaces via software-based interfaces call Bridged Virtual Interfaces (BVI).

vLANs can be extended through routers by mapping the vLANs to bridge-groups. IRB can be used to perform inter-vLAN routing between the extended vLANs and to interconnect the extended vLANs to the routed domain inside the same Cisco IOS based router.

Figure 6. Integrated Routing and Bridging

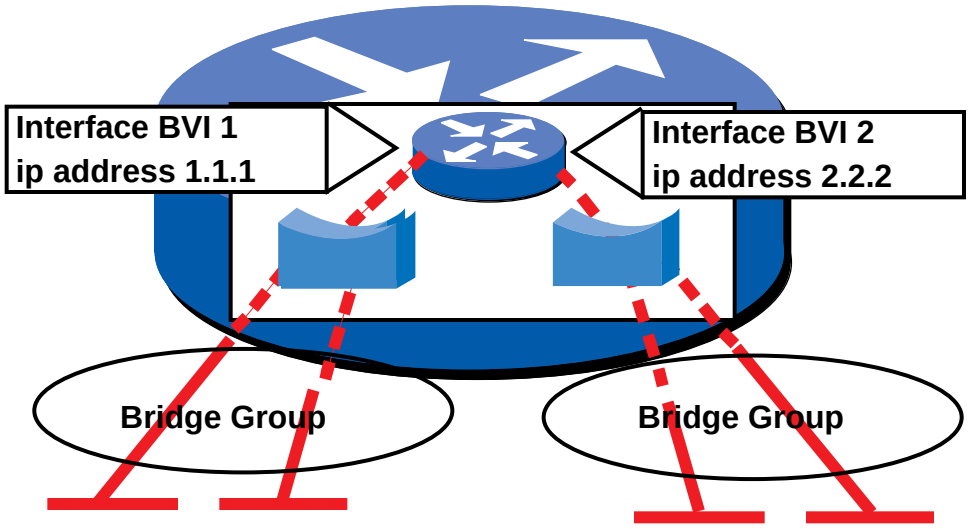

Benefits: The benefits of Integrated Routing and Bridging include:

- **Scalable, efficient integration of layer 2 and layer 3 domains** -- The IRB functionality allows you to extend the bridge domains or vLANs across routers while maintaining the ability to interconnect them to the routed domains through the same router.
- **Layer 3 address conservation** -- You can extend the bridge domains and the vLAN environments across the routers to conserve the layer 3 address space and still use the same router to interconnect the vLANs and bridged domains to the routed domain.
- **Flexible network reconfiguration** -- Network administrators gain the flexibility of being able to extend the bridge domain across the router's interfaces to provide temporary solution for moves, adds, and changes. This can be useful during migration from a bridged environment to a routed environment, or when making address changes on a scheduled basis.

Considerations:

- Currently, IRB supports three protocols: IP, IPX, and AppleTalk, in both fast switching and process switching modes.
- IRB is not supported on cbus platforms (AGS+ and 7000).
- IRB is supported for transparent bridging, but not for source route bridging.
- IRB is supported on all media-type interfaces except X.25 and ISDN bridged interfaces.
- IRB and Concurrent Routing and Bridging (CRB) cannot operate at the same time.

Product Marketing Contact: Amir Khan

2 Desktop Protocols

2.1 AppleTalk Features

2.1.1 AppleTalk Load Balancing

Description: This feature allows AppleTalk data traffic to be distributed more evenly across redundant links in a network.

Benefits: AppleTalk load balancing can reduce network costs by allowing more efficient use of network resources. Network reliability is improved because the chance that network paths between nodes will become overloaded is reduced. For convenience, load balancing is provided for networks using native AppleTalk routing protocols such as Routing Table Maintenance Protocol (RTMP) and Enhanced IGRP.

Considerations: AppleTalk load balancing operates with process and fast switching.

Product Marketing Contact: Roger Farnsworth

2.2 Novell Features

2.2.1 Display SAP by Name

Description: This feature allows network managers to display Service Advertisement Protocol (SAP) entries that match a particular server name or other specific value. The current command that displays IPX servers has been extended to allow the use of any regular expression (including supported special characters) for matching against the router's SAP table.

Benefits: The job of supporting ever larger Novell IPX networks is a constant challenge to today's network managers. Cisco continues to deliver tools to ease the administrative burden associated with these growing internetworks. By providing network managers with the ability to more effectively display SAP table entries, network diagnostic and maintenance costs can be reduced.

Considerations: None

Product Marketing Contact: Roger Farnsworth

2.2.2 IPX Access Control List Violation Logging

Description: With this feature, routers can use existing router logging facilities to log IPX access control list (ACL) violations whenever a packet matches a particular access-list entry. The first packet to match an entry is logged immediately; updates are sent at approximately five-minute intervals.

This feature allows logging of:

- Source and destination addresses
- Source and destination socket numbers
- Protocol (or packet) type (for example, IPX, SPX, or NCP)
- Action taken (permit/deny)

Benefits: IPX ACL violation logging provides network managers with additional functionality to more effectively control access to important network resources. Along with the responsibility for controlling the flow and accessibility of information in today's networks, network managers are increasingly challenged with detecting and correcting breaches of network security. IPX ACL violation logging provides a method of establishing a centralized reporting system in order to provide early notification of attempted unauthorized access.

Considerations: Matching packets and logging-enabled ACLs are sent at the process level. Router logging facilities use the IP protocol.

Product Marketing Contact: Roger Farnsworth

2.2.3 Plain English IPX Access List

Description: Through the use of this feature, the most common protocol and socket numbers used in IPX extended access control lists (ACLs) can be specified by either name or number (as opposed to just the cryptic numbers required in the past.)

Protocol types supported include RIP, SAP, NCP, and NetBIOS. Supported socket types include Novell Diagnostics Packet Enhanced IGRP, and NLSP.

Benefits: Plain English IPX Access Lists greatly reduce the complexity and readability of IPX extended access control lists, reducing network management expense by making it easier to build and analyze the access control mechanisms used in IPX networks. By reducing the complexity of the commands, it is easier than ever to control the security and accessibility of IPX network resources.

Considerations: None

Product Marketing Contact: Roger Farnsworth

3 Wide-Area Networking Features

3.1 ISDN/DDR Enhancements

3.1.1 Multichassis Multilink PPP (MMP)

Description: Multichassis Multilink Point-to-Point Protocol (MMP) extends Multilink PPP (MP) by providing a mechanism to aggregate B-channels transparently across multiple routers or access servers. MMP defines the methodology for sharing individual links in a MP bundle across multiple, independent platforms. The primary application for MMP is the ISDN dial-up pool; however, it can also be used in a mixed technology environment.

MMP is based on the concept of a "stackgroup" -- a group of routers or access servers that operate as a group when receiving MP calls. Any member of the stackgroup can answer any call into the single access number applied to all WAN interfaces. Typically, the access number corresponds to a telco hunt group.

Cross platform aggregation is performed via tunneling between members of a stackgroup using the Layer 2 Forwarding (L2F) protocol, a draft IETF standard. This protocol is also the basis of Cisco virtual private dial-up network offering.

MMP is flexible and scalable. Because the L2F protocol is IP-based, members of a stackgroup can be connected over many types of LAN or WAN media. Stackgroup size can be increased by increasing the bandwidth available to the L2F protocol -- for example, by moving from shared to switched ethernet.

Figure 7. Multichassis Multilink PPP

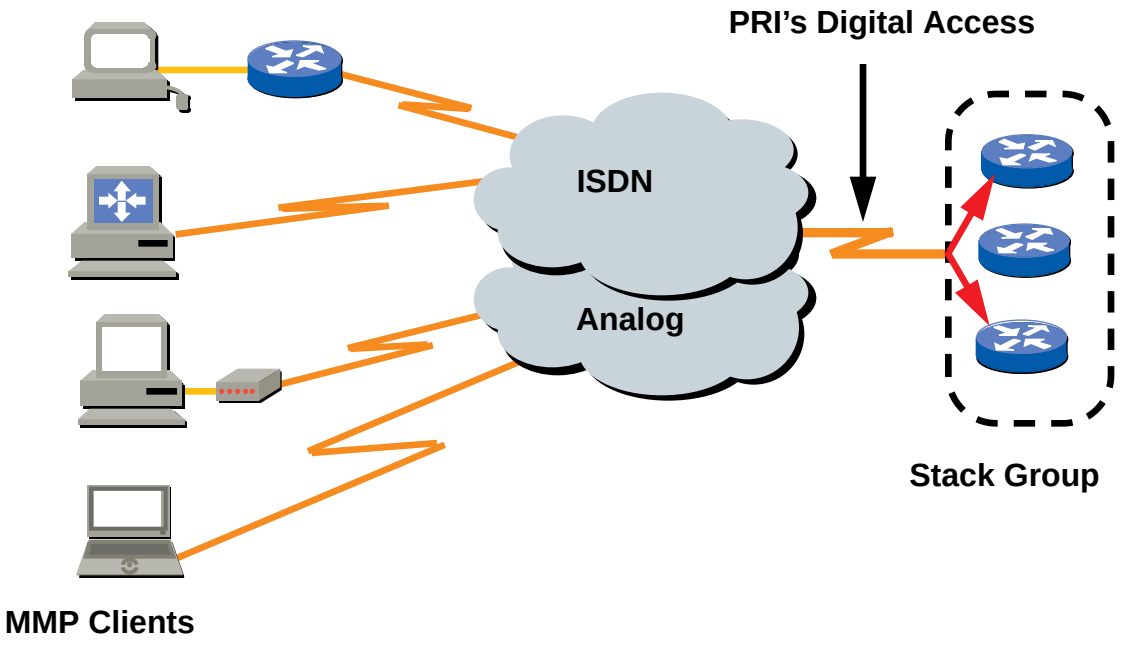

Benefits: Multichassis Multilink PPP provides a number of significant benefits over existing dial-up pool solutions. These include:

- Organizations can scale their core network as needed. With MMP, new devices can be added to the dial-up pool at any time.
- MMP is less CPU intensive than MP. The load for reassembly and resequencing can be shared across all devices in the stackgroup. This allows for the use of less expensive platforms in large dial-up pools.
- MMP provides an interoperable multivendor solution since it does not require any special software capabilities at the remote sites. The only remote requirement is support for industry standard MP (RFC 1717).

Considerations: MMP is a versatile, platform independent technology allowing stackgroups to be formed from differing platform types such as the Cisco 2500, Cisco 4000 and Cisco 7000 series routers.

Universal access servers such as the Cisco 5200 should not be combined with ISDN-only access servers such as the Cisco 4000 series router in a MMP stackgroup. Since calls are allocated by the Central Office in an arbitrary manner it is possible that this scenario could lead to an analogue call being delivered to a digital-only access server.

Product Marketing Contact: Kevin Dickson

3.1.2 Virtual Private Dial-up Network

Description: Virtual Private Dial-up Network (VPDN) functionality is based on the Layer 2 Forwarding (L2F) specification which Cisco has proposed as an industry standard to the Internet Engineering Task Force (IETF). VPDN allows users from multiple disparate domains to gain secure access to their corporate home gateways via public networks or the Internet.

Service providers who wish to offer private dial-up network services can use VPDN to provide a single telephone number for all their client organizations. A customer can use dial-up access to a local point-of presence where the access server identifies the customer by PPP user name. The PPP user name is also used to establish a home gateway destination. Once the home

gateway is identified, the access server builds a secure tunnel across the service provider's backbone to the customer's home gateway. The PPP session is also transported to this home gateway, where local security measures can ensure the person is allowed access to the network behind the home gateway.

Of special interest to service providers is VPDN's independence of WAN technology. Since L2F is TCP/IP-based, it can be used over any type of service provider backbone network.

Benefits: VPDN offers many benefits to the service provider and service provider customer. These include:-

- No special software is required by the end-user or his home gateway.
- Standard authentication techniques such as TACACS or RADIUS can be employed at the home gateway to ensure that security is maintained. This authentication is independent of anything done by the service provider.
- Address assignment is handled by the users home gateway. No service provider addressing is used.

Considerations: None.

Product Marketing Contact: Kevin Dickson

3.1.3 Dialer Profiles

Description: Dialer Profiles is a new dialer feature that allows the user to separate the "logical" portion of the configuration (network layer, encapsulation, dialer parameters) from that of the interface used to place or receive calls.

Dialer Profiles are not true sub-interfaces in that the Central Office switch still determines which connection is used. Incoming calls are tied to a specific profile, based on either the Caller ID or the PPP user name of the calling party.

Figure 8. Dialer Profiles.

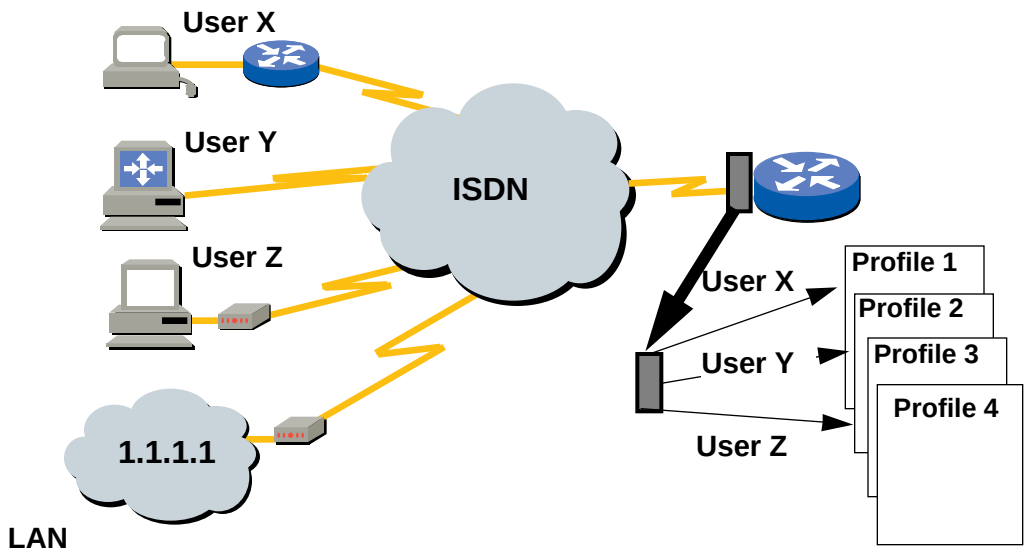

Benefits: Dialer Profile extends the flexibility of current dial-up configurations. For example, on a single ISDN PRI or PRI rotary group it is now possible to allocate separate profiles for different classes of user. These profiles may define normal DDR usage or backup usage.

Considerations: Each dialer profile uses an Interface Descripter Block (IDB) distinct from the IDB of the physical interface used to place or receive calls. When a call is established, both IDBs are bound together so that traffic can flow. As a result, Dialer Profiles uses more IDBs than normal DDR.

This initial release of dialer profiles does not support Frame Relay, X.25, or LAPB encapsulation on DDR links or Snapshot Routing capabilities.

Product Marketing Contact: Kevin Dickson

3.1.4 Combinet Packet Protocol (CPP) Support

Description: Combinet Packet Protocol (CPP) is a proprietary encapsulation used by legacy Combinet products for data transport. CPP also defines a methodology for performing compression and load sharing across ISDN links. The Cisco IOS software implementation of CPP supports both compression and load sharing using this proprietary encapsulation.

Benefits: A large installed base of early Combinet product users cannot upgrade to later software releases that support interoperability standards such as PPP. With CPP support, these users can integrate their existing product base into new Cisco IOS-based internetworks.

Considerations: CPP does not provide many of the functions available in Cisco's implementation of the PPP standards. These functions include address negotiation and support for protocols like AppleTalk. Where possible, customers should migrate to software that supports PPP.

Product Marketing Contact: Kevin Dickson

3.1.5 Half Bridge/Half Router for Combinet Packet Protocol (CPP) and PPP

Description: Half bridge/half router allows low-end, simply configured bridge devices to bridge either PPP or Combinet Packet Protocol (CPP) encapsulated data to a Cisco IOS core router. Half bridge/half router is designed for networks that have a small remote Ethernet segments with a single PPP- or CPP-compatible bridging device connected to a core network. The serial or ISDN interface on the core router appears as a virtual Ethernet port to the network. Layer 3 data packets transported across this type of link are first encapsulated within an Ethernet encapsulation. A PPP or CPP bridging header is then added. This facility allows bridged traffic arriving at the core device to be routed from that point on.

Benefits: As remote access devices become more prevalent, there is a genuine requirement to greatly simplify the initial configuration of these units. Half bridge offers the simplicity of a bridging configuration at the remote site coupled with the benefits of a routing configuration at the core site.

Considerations: This feature is process switched.

Product Marketing Contact: Kevin Dickson

3.2 Frame Relay Enhancements

3.2.1 Frame Relay SVC Support (DTE)

Description: Currently, access to Frame Relay networks is through private leased lines at speeds ranging from 56 kbps to 45 Mbps. Bandwidth within the Frame Relay network is permanently committed to providing Permanent Virtual Circuits (PVCs) between the endpoints. Switched Virtual Circuits (SVCs) allow access through a Frame Relay network by setting up a path to the destination endpoints only when the need arises. This is similar to X.25 SVCs which allow connections to be set up and torn down based upon data traffic requirements. Although SVCs entail link set up and tear down overhead, the benefit is that the VC is only established when data must be transferred. Therefore, the number of VCs is proportional to the number of actual conversations between sites rather than the number of sites.

Figure 9. Frame Relay SVC Support (DTE)

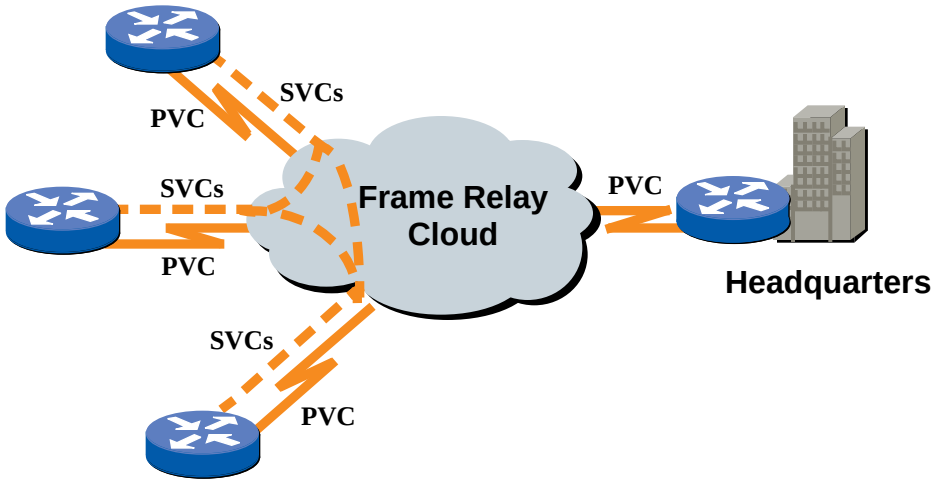

One PVC from each branch office to headquarters.

One SVC from each branch office to every other office.

Benefits: Frame Relay SVCs offer the following benefits:

- Cost savings via usage-based pricing instead of fixed pricing for a PVC connection. This is similar to the billing system used for standard telephone service.
- Dynamic modification of network topologies with any-to-any connectivity.
- Dynamic network bandwidth allocation or bandwidth-on-demand for large data transfers such as FTP traffic.
- Backup for PVC backbones.
- Conservation of resources in private networks. Connetions and router CPU are allocated only when the connection is required to transfer data.

Considerations: The following facilities are necessary before Frame Relay SVCs can operate:

- Frame Relay SVC must be supported by the Frame Relay switches used in the network.
- A Physical Local Loop Connection, such as a leased or dedicated line, must exist between the router (DTE) and the local Frame Relay switch.

Product Marketing Contact: Sanjay Bhardwaj

3.2.2 Traffic Shaping over Frame Relay

Description: The Frame Relay protocol defines several parameters that are useful for managing network traffic congestion. These include: Committed Information Rate (CIR), Forward/Backward Explicit Congestion Notification (FECN/BECN), and Discard Eligibility (DE) bit. Cisco already provides support for: FECN for DECnet and OSI, BECN for SNA traffic using direct LLC2 encapsulation via RFC 1490, and DE bit support. The Frame Relay Traffic Shaping feature builds upon this support by providing the following three capabilities:

- **Rate Enforcement on a per Virtual Circuit (VC) basis** -- A peak rate can be configured to limit outbound traffic to either the CIR or some other defined value such as the Excess Information Rate (EIR).
- **Generalized BECN support on a per VC basis** -- The router can monitor BECNs and throttle traffic based upon BECN marked packet feedback from the Frame Relay network.
- **Priority/Custom/Weighted-Fair Queuing (PQ/CQ/WFQ) support at the VC level** -- This allows for finer granularity in the prioritization and queuing of traffic, providing more control over the traffic flow on an individual VC.

These features improve the scalability and performance of a Frame Relay network by increasing the density of VCs and improving response time. The Frame Relay Traffic Shaping feature applies to Frame Relay PVCs and SVCs.

Figure 10. Traffic Shaping over Frame Relay

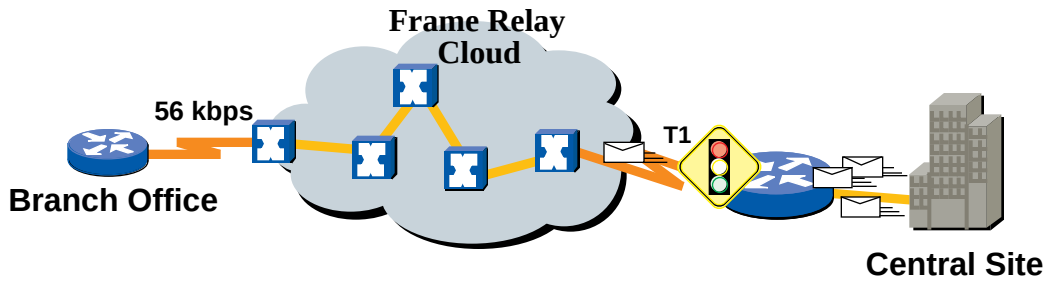

Benefits: The Frame Relay Traffic Shaping feature provides the following benefits:

- Eliminates bottlenecks in Frame Relay network topologies with high speed connections at the central site, and low speed connections at the branch sites. Rate Enforcement can be used to limit the rate at which data is sent on the VC at the central site. Rate Enforcement can also be used in conjunction with the existing DLCI Prioritization feature to further improve performance in this situation.
- Provides a mechanism for sharing media by multiple VCs. Rate Enforcement allows the transmission speed used by the router to be controlled by criteria other than line speed, such as the CIR or EIR. The Rate Enforcement feature can also be used to pre-allocate bandwidth to each VC, creating a Virtual Time Division Multiplexing network.
- Dynamically throttles traffic, based on information contained in BECN-tagged packets received from the network. With BECN based throttling, packets are held in the router's buffers to reduce the data flow from the router into the Frame Relay network. The throttling is done on a per VC basis and the transmission rate is adjusted based on the number of BECN-tagged packets received.
- Defines queuing at the VC or sub-interface level. Custom Queuing with the Per VC Queuing and Rate Enforcement capabilities enable Frame Relay VCs to be configured to carry multiple traffic types (such as IP, SNA and IPX), with bandwidth guaranteed for each traffic type.

Considerations: The three components of the Traffic Shaping for Frame Relay feature require the router to buffer packets to control traffic flow and compute data rate tables. Because of this router memory and CPU utilization, these features must be used judiciously to regulate critical traffic flows while not degrading overall Frame Relay performance.

Product Marketing Contact: Sanjay Bhardwaj

3.3 ATM Enhancements

3.3.1 Simple Server Redundancy Protocol (SSRP) for LAN Emulation

Description: The Simple Server Redundancy Protocol (SSRP) provides stand-by redundancy for the following services used by clients in an ATM LAN Emulation (LAN) network: LAN Emulation Configuration Server (LECS), LAN Emulation Server (LES), and Broadcast and Unknown Server (BUS). As many as four LECS can be defined (triple redundancy) in this release. LECS addresses can be defined in ILMI on a per port basis in the LightStream 1010 and the LightStream 100 switches.

Benefits: LAN Emulation uses one LES/BUS per emulated LAN and one LECS per multiple emulated LANs. These service components represent single points of failure for each emulated LAN. SSRP removes these single points of failure, providing network managers the redundancy they need for campus ATM backbones with LAN Emulation without adding administrative overhead. A completely redundant, dual-homed ATM backbone can be built without any failure points when SSRP is combined with Hot Standby Router Protocol (HSRP), the dual-phy LANE card for the Catalyst 5000, and support for Spanning Tree on a per VLAN-basis.

Considerations: Full implementation of SSRP requires Cisco platforms. Currently, LECS and LES/BUS are available on the Cisco 7000 series, Cisco 7500 series, Cisco 4000 series routers, and the Catalyst 5000. Any standard LAN Emulation 1.0 Client (LEC) can take advantage of the LECS and LES/BUS redundancy in SSRP without additional capability. A LEC can take advantage of this mechanism in SSRP to find a backup LECS as follows :

- Supporting standard ILMI completely, and
- Contacting successive LECs should one not respond during initialization.

The Catalyst 5000 LAN Emulation module will support SSRP with release 3.1.

Product Marketing Contact: David Benham

3.3.2 Hot Standby Router Protocol (HSRP) support for LAN Emulation

Description: Most hosts are configured with, or discover, their default gateway (router). If there is more than one router connected to an emulated LAN, Cisco's Hot Stand-by Router Protocol (HSRP) allows one of those routers to monitor the status of the other and take over the functions of that router should it fail or become unavailable.

Benefits: HSRP provides inter-ELAN (or inter-VLAN) routing redundancy. HSRP over LANE is transparent to hosts expecting to always to be able to reach their default gateway (router). Without HSRP, IP hosts would need to be configured with RIP to recover from a failure of its default gateway. This method can result in a 10 minute delay before the host can use its second default gateway. A completely redundant, dual-homed ATM backbone can be built without any failure points when HSRP is combined with Simple Server Redundancy Protocol (SSRP), the dual-phy LANE card for the Catalyst 5000, and support for Spanning Tree on a per VLAN-basis.

Considerations: HSRP is a unique protocol developed by Cisco and used only by Cisco IOS software-based routers. HSRP over LAN Emulation will be available on the Cisco 7000 series, Cisco 7500 series and Cisco 4000 series routers, which all have an ATM interface. Any LAN Emulation Client (LEC) will benefit from HSRP, as HSRP is transparent to the LEC. Cisco IOS software-based routers have the additional benefit that they can fully participate in HSRP, acting as a back up gateway.

Product Marketing Contact: David Benham

3.3.3 DECnet and VINES Routing support for LAN Emulation

Description: This adds the ability to route DECnet and VINES from a sub-interface on an ATM router port running LAN Emulation to any other sub-interface on an ATM router port or any other router port. This joins the current capability to route IP, IPX, and AppleTalk over LAN Emulation sub-interfaces.

Benefits: Support for DECnet and VINES routing between VLANs for ATM LAN Emulation

Considerations: DECnet Phase IV

Product Marketing Contact: David Benham

3.3.4 UNI 3.1 Signaling Support

Description: Cisco IOS software Release 11.2 supports UNI 3.1 signaling. The ATM Forum submitted the UNI 3.0 signaling specification to the ITU, which subsequently made changes to the SSCOP encapsulation used to make signaling reliable. UNI 3.1 was published later by the ATM Forum to align with the ITU, otherwise there is no difference in functionality between UNI 3.0, currently supported on all Cisco ATM platforms, and UNI 3.1.

Benefits: The full breadth of UNI signaling protocol support is available.

Considerations: none.

Product Marketing Contact: David Benham

3.3.5 Rate Queues for SVCs per sub-interface

Description: Currently, only PVCs can be assigned to a particular rate queue (which defines traffic shaping parameters), whereas SVCs always fall into the default rate queue. Now, all SVCs in a sub-interface (such as a sub-interface configured to run RFC 1577 Classical IP over ATM) are put into the rate queue assigned to that sub-interface.

Benefits: All connections (not just PVCs) in a sub-interface can be traffic shaped identically to match a service contract, for example with a service provider.

Considerations: PVCs on any sub-interface can still be assigned to any rate queue, so if this isn't desired, be careful with configuration commands.

Product Marketing Contact: David Benham

3.3.6 AToM MIB Support

Description: This provides support for the AToM MIB, described in IETF RFC 1695, which defines configuration information as well as error and cell-level counters. Cisco IOS software Release 11.2 provides a standard AToM MIB instrumentation for many of the counters already provided in the router's ATM interfaces.

Benefits: AToM MIB instrumentation is used by network management applications, such as Cisco's AtmDirector, to perform topology auto-discovery and status checking.

Considerations: None

Product Marketing Contact: David Benham

3.4 Core Enhancements

3.4.1 NetFlow Switching

Description: NetFlow Switching is a new software switching mechanism that allows Cisco routers to combine high-performance network-layer switching with the application of network services. To achieve this high performance, NetFlow Switching identifies traffic flows between internetwork endpoints and then, on a connection-oriented basis, switches packets in these streams at the same time that it applies relevant services. By identifying flows using both network-layer and transport-layer information, NetFlow Switching allows Cisco IOS services to be applied on a per-user, per-application basis.

Benefits: With NetFlow Switching, network users can extend their use of existing Cisco IOS services, such as security access lists or the collection of traffic statistics, without paying the performance penalty usually associated with such processing-intensive functions. This increase in performance allows these services to be used in more places within the network and on a larger scale. Extending network security is increasingly important as networks need to support access from remote

users and across public Internet services. Detailed information on traffic flows helps network managers to grow their networks in the most cost-effective way. With NetFlow Switching, network administrators finally have “call detail recording” information for their data networks.

NetFlow Switching provides increased performance for the application of existing Cisco IOS services such as security access lists and accounting. Previously, system performance could be affected by as much as 30 percent for each service invoked. With NetFlow Switching, system switching performance can be maintained within 10 to 15 percent of optimum levels -- for all supported services. As with any connection-oriented technique, the performance of NetFlow Switching is affected by the total number of active flows.

Considerations: Cisco’s initial implementation of NetFlow Switching supports Internet Protocol (IP) traffic over all interface types and provides optimal performance with Ethernet, Fiber Distributed Data Interface (FDDI), and High-Level Data Link Control (HDLC) serial interfaces. Cisco will extend NetFlow Switching to handle Internetwork Packet Exchange (IPX) traffic later in 1996.

NetFlow Switching is supported onthe Cisco 7500 series and Cisco 7000 series routers with a Route/Switch Processor (RSP). On these routers, NetFlow Switching can operate on the master RSP or on a distributed basis on individual Versatile Interface Processors (VIPs).

Product Marketing Contact: Steve Collen

4 IBM Functionality

4.1 APPN Enhancements

4.1.1 APPN Central Resource Registration

Description: APPN Central Resource Registration (CRR) support allows a Cisco IOS software-based router acting as a network node (NN) to register the resources of end nodes (ENs) to the Central Directory Service (CDS) on Advanced Communication Facility/Virtual Telecommunication Access Method (ACF/VTAM). A Cisco IOS NN will now register resource names with a VTAN CDS as soon as it establishes connectivity with it. Prior to this enhancement, the router acting as a NN could not register EN resources. ACF/VTAM could, however, query the router to find these resources.

Figure 11. APPN Central Resource Registration

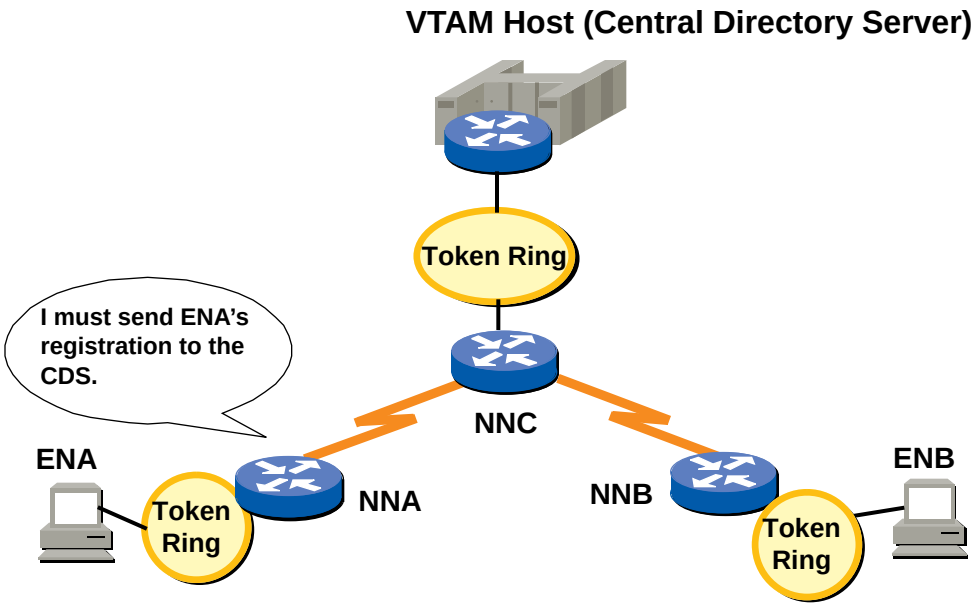

Benefits: The CDS reduces broadcast traffic in the network. Without an active CDS on ACF/VTAM, the NN must send a broadcast message to the network to locate non-local resources required for a session. With an active CDS, the NN sends a single request directly to the CDS for the location of the resource. A network broadcast is used only if the resource has not registered with the CDS.

Considerations: ACF/VTAM must be configured as a CDS. The Cisco IOS NN learns of the capability when network topology is exchanged. To most effectively use the CDS, ENs should register the resources with the NN. Depending on the EN implementation, registration may occur automatically, may require configuration on the EN, or may not be a function of the EN.

Product Marketing Contact: Betsy Huber

4.1.2 APPN DLUR MIB

Description: The existing APPN Management Information Base (MIB) does not contain information about Dependent Logical Units (DLUs) accessing the APPN network through the DLU Requester (DLUR) function in the Cisco IOS NN. A standard MIB for DLUR has been defined by the APPN Implementers Workshop (AIW), the standards body for APPN, and is implemented in this release of the Cisco IOS software.

Benefits: With the APPN DLUR MIB, users have access to information collected about the DLUR function in the Cisco IOS NN and the DLUs attached to it for more complete network management information.

Considerations: The availability of the MIB is independent of the management application. Enhancements to CiscoWorks to exploit the information contained in the APPN DLUR MIB will be announced separately.

Product Marketing Contact: Betsy Huber

4.2 Native Client Interface Architecture (NCIA) Server

Description: The Native Client Interface Architecture (NCIA) architecture, introduced by Cisco Systems for access of IBM SNA applications over routed internetworks, has been enhanced to be more flexible and scalable. The NCIA Client, implemented in the client workstation, encapsulates the full SNA stack inside TCP/IP packets. These packets are sent to the NCIA Server implemented in Cisco IOS software. The NCIA Server de-encapsulates the TCP/IP packet and sends the LLC data to the host processor via RSRB or DLSw+.

Figure 12. Native Client Interface Architecture Server.

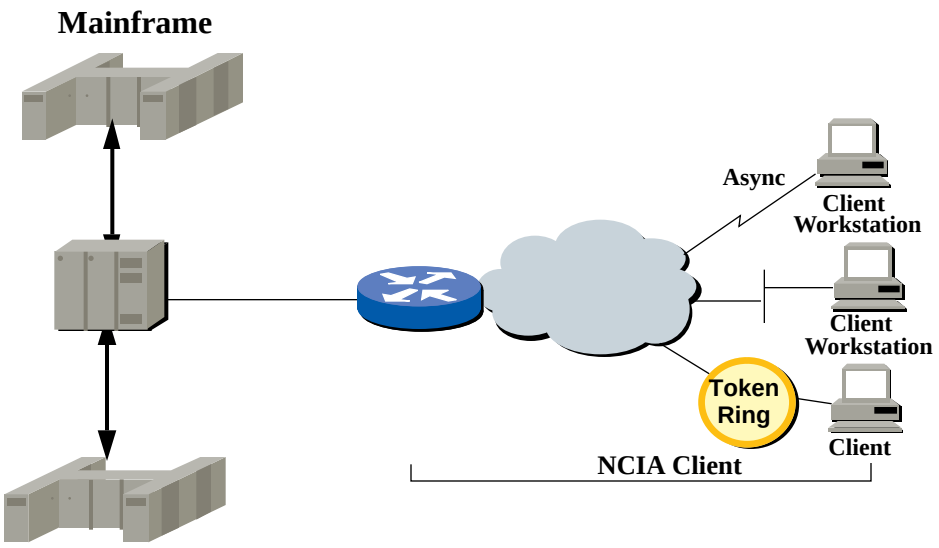

Benefits: The NCIA Server supports SNA and NetBIOS sessions over a variety of LAN and WAN connections, including dial-up connections. The NCIA architecture supports clients with full SNA stacks -- providing all advanced SNA capabilities, unlike some split-stack solutions. The new NCIA Server enhancements provide the following additional benefits:

- Client configuration is simplified. It is no longer necessary to predefine ring numbers, and the NCIA Server supports optional dynamic assignment of MAC addresses. There is no Logical Link Control, type 2 (LLC2), at the client. The client is configured as an end station, not a router peer.
- The solution is more scalable. The limit is based on the number of LLC connections in the central site router rather than RSRB peer connections.
- The protocol used between client and server is efficient.

Considerations: Each client is a full SNA PU with one or more LUs. As such, each device requires one LLC connection at the central site router. The Cisco 4700 currently supports 3000-4000 LLC connections.

Product Marketing Contact: Lisa Lindgren

4.3 TN3270 Server

Description: The TN3270 Server is a new feature of the Channel Interface Processor (CIP) of the Cisco 7000 family of routers. The TN3270 Server allows TN3270 clients access to IBM and IBM-compatible mainframes without the limitations of existing alternatives. It off-loads 100 percent of the TCP/IP and TN3270 cycles from the mainframe, and offers a robust, scalable and dynamic implementation that meets the stringent requirements of the Data Center.

Benefits: The TN3270 Server on the CIP is an extremely scalable solution when compared to other alternatives on the market. It supports up to 8000 concurrent sessions, while most external gateway solutions can only support up to 1000-2000 sessions. The TN3270 Server offers the following advanced capabilities:

- **Load Balancing and Redundancy** -- Provides effective utilization of CIP resources and more consistent response times.
- **End-to-End Session Visibility** -- Provides enhanced management of resources.
- **SNA Session Switching** -- Off-loads VTAM by providing session routing.
- **TN3270E Support** -- In combination with a TN3270E client, provides advanced SNA management and SNA functionality, cluding printer support.
- **Dynamic Definition of Dependent LUs** -- Provides simplified configuration and network definition at the router and in VTAM.
- **Dynamic Allocation of LUs** -- Makes efficient use of LU pool resources while supporting multiple SNA model types.

Considerations: TN3270 Server requires 32 MB of CIP DRAM to support up to 4000 sessions, and 64 MB to support 8000 sessions. TN3270 Server can run concurrently with any of the other CIP applications (IP Datagram, TCP/IP Off-load, or CSNA), but operation of any of these features will affect the total number of sessions supported due to contention for CIP processor cycles.

Product Marketing Contact: Lisa Lindgren

4.4 Fast Switched Source-Route Translational Bridging (SR/TLB)

Description: With Cisco IOS software release 11.2, SR/TLB is Fast Switched. No queuing is done and resource utilization is low. This enhancement is on by default, but can be disabled. It is supported across all router platforms.

Benefits: Fast Switched SR/TLB improves performance on all platforms by a factor of at least 2, and for the Cisco 4500 and Cisco 4700, by a factor of 3. It is ideal for IBM environments (for example, where low-cost Ethernet adapters are being installed on campus, but Token Ring connectivity to a FEP is still required), and for campus environments with a mix of Token Ring and Ethernet LANs and/or switches that rely on the Cisco IOS software for translational bridging.

Considerations: none.

Product Marketing Contact: Donna Kidder

4.5 Data Link Switching+ (DLSw+) Features and Enhancements

With Cisco IOS software, release 11.2, the following features are supported:

- LAN Network Manager (LNM) over DLSw+
- Native Service Point (NSP) over DLSw+
- Down Stream Physical Unit (DSPU) over DLSw+
- Advanced Peer-to-Peer Networking (APPN) over DLSw+
- SRB over FDDI to DLSw+

These features had previously been available with Remote Source-Route Bridging (RSRB). To provide these features, the Cisco IOS software uses a component known as Virtual Data Link Control (VDLC) that allows one software component to use another software component as a data link.

4.5.1 LAN Network Manager (LNM) over DLSw+

Description: LAN Network Manager (LNM) over DLSw+ allows DLSw+ to be used in Token Ring networks that are managed via IBM's LNM software.

Benefit: Using this feature, LNM can be used to manage Token Ring LANs, Control Access Units (CAUs), and Token Ring attached devices over a DLSw+ network. All management functions continue to operate as they would in an RSRB network or source route bridged network.

Considerations: None.

Product Marketing Contact: Donna Kidder

4.5.2 Native Service Point (NSP) over DLSw+

Description: Native Service Point (NSP) over DLSw+ allows Cisco's NSP feature to be used in conjunction with DLSw+ in the same router.

Benefit: Using this feature, NSP can be configured in remote routers, and DLSw+ can provide the path for the remote service point PU to communicate with NetView. This allows full management visibility of resources from a NetView 390 console, while concurrently offering the value-added features of DLSw+ in an SNA network.

Considerations: None.

Product Marketing Contact: Donna Kidder

4.5.3 Down Stream Physical Unit (DSPU) over DLSw+

Description: Down Stream Physical Unit (DSPU) over DLSw+ allows Cisco's DSPU feature to operate in conjunction with DLSw+ in the same router. DLSw+ can be used either upstream (towards the mainframe) or downstream (away from the mainframe) of DSPU.

Benefit: DSPU concentration consolidates the appearance of up to 255 physical units into a single PU appearance to VTAM, minimizing memory and cycles in central site resources (VTAM, NCP, and routers) and speeding network startup. Used in conjunction with DLSw+, network availability and scalability can be maximized.

Considerations: None.

Product Marketing Contact: Donna Kidder

4.5.4 Advanced Peer-to-Peer Networking (APPN) over DLSw+

Description: Advanced Peer-to-Peer Networking (APPN) over DLSw+ allows Cisco's APPN feature to be used in conjunction with DLSw+ in the same router.

Benefit: With this feature, DLSw+ can be used as a low cost way to access an APPN backbone or APPN in the data center. In addition, DLSw+ can be used as a transport for APPN, providing non-disruptive recovery from failures and high speed intermediate routing. In this case, the DLSw+ network appears as a connection network to the APPN network nodes (NNs).

Considerations: None.

Product Marketing Contact: Donna Kidder

4.5.5 Source Route Bridging (SRB) over FDDI to DLSw+

Description: This feature allows access to DLSw+ over source route bridged FDDI LANs. In the past, the supported local DLCs were only Token Ring, Ethernet, or SDLC. With this extension, Token Ring-attached devices can access a DLSw+ router using source route bridging over an FDDI backbone. At the remote site, the device can be attached over Token Ring, Ethernet, SDLC, or FDDI. This is useful either in environments with Token Ring switches that use FDDI as a campus backbone or environments with Cisco 7000 and Cisco 7500 series routers providing SRB over an FDDI backbone.

Benefit: Allows SRB over FDDI to provide the highest speed access between campus resources, while concurrently allowing DLSw+ for access to remote resources.

Consideration: Currently, SRB over FDDI is supported by the Cisco 7000 and Cisco 7500 series platforms only.

Product Marketing Contact: Donna Kidder

4.6 Response Time Reporter

Description: The Response Time Reporter (RTR) feature allows you to monitor network performance, network resources and applications by measuring response times and availability. RTR statistics can be used to perform troubleshooting, problem notifications and pre-problem analysis. RTR offers enhanced functionality over a similar IBM product, NetView Performance Monitor.

Figure 13. Respose Time Reporter

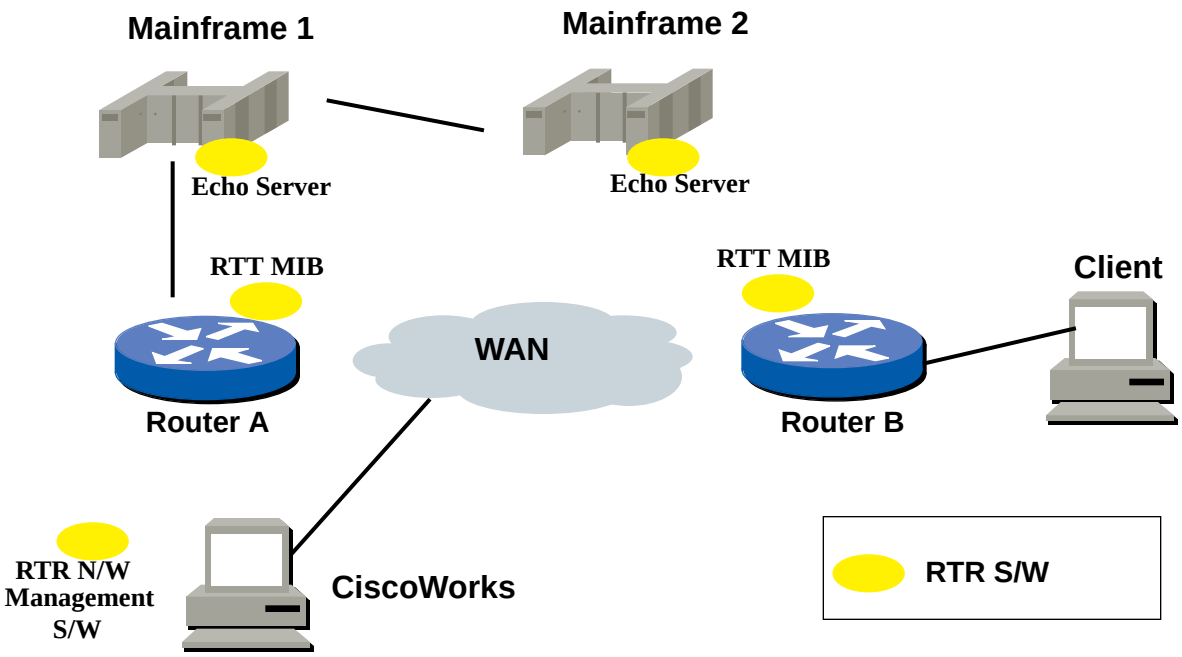

Benefits: RTR enables the following functions to be performed:

- Troubleshoot problems by checking the time delays between devices (such as a router and a MVS host) and the time delays on the path from the source device to the destination device at the protocol level.
- Send SNMP traps and/or SNA Alerts/Resolutions when one of the following has occurred: a user-configured threshold is exceeded, a connection is lost and reestablished, or a timeout occurs and clears. Thresholds can also be used to trigger additional collection of time delay statistics.

- Perform pre-problem analysis by scheduling the RTR and collecting the results as history and accumulated statistics. The statistics can be used to model and predict future network topologies.

Considerations: The RTR feature is currently available only with the Cisco IBM feature set. A CiscoWorks Blue network management application will be available to support the RTR feature. Both the CiscoWorks Blue network management application and the router use the Cisco Round Trip Time Monitor (RTTMON) MIB. This MIB is also available with release 11.2 of Cisco IOS software.

Product Marketing Contact: Lori Bush

5 Security Features

5.1 Router Authentication and Network-Layer Encryption

Description: This feature provides a mechanism for secure data transmission. It consists of two components:

- **Router Authentication** -- Prior to passing encrypted traffic, two routers perform a one time, two-way authentication by exchanging Digital Signature Standard (DSS) public keys. The hash signatures of these keys are compared to authenticate the routers.
- **Network-Layer Encryption** -- For IP payload encryption, the routers use Diffie-Hellman key exchange to securely generate a DES 40- or 56- bit session key. New session keys are generated on a configurable basis. Encryption policy is set by "crypto-maps" that use extended IP Access Lists to define which network, subnet, host, or protocol pairs are to be encrypted between routers.

Figure 14. Router Authentication and Network-Layer Encryption

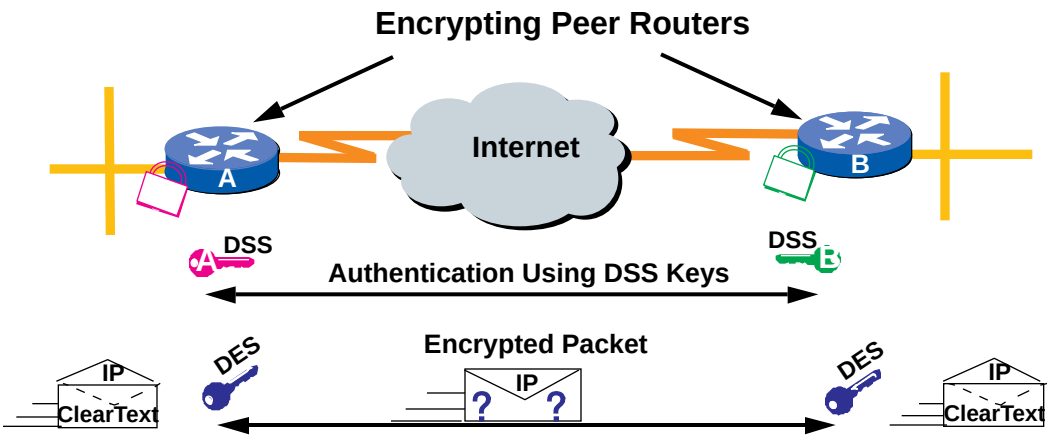

Benefits: Customers can now run high assurance, confidential connections over public or untrusted IP networks. This feature can be used to build multiprotocol Virtual Private Networks (VPNs), using encrypted Generic Routing Encapsulation (GRE) tunnels. It can also be used to deploy secure telecommuting services, Intranet privacy, and virtual collaborative or community-of-interest networks.

Considerations: All components of this feature are subject to International Traffic in Arms Regulations (ITAR) export restrictions. Encryption is currently IP only, though it does support multiprotocol GRE tunnels. This feature is most appropriately deployed in a relatively small number of routers, with a logically flat or star-shaped encryption topology.

Load-sharing of the encryption/decryption function is not supported. Without a Certification Authority (CA), the one-time authentication effort increases exponentially with the number of routers. Router authentication requires the network administrator to compare the hashes produced by the routers. This version of encryption is not IPSEC compliant.

Product Marketing Contact: Elizabeth Kaufman

5.2 TACACS+ Enhancements

5.2.1 TACACS+ Single Connection

Description: Single Connection is an enhancement to the Network Access Server (NAS) that increases the number of transactions per second supported. Prior to this enhancements, separate TCP connections would be opened and closed for each of the TACACS+ services: authentication, authorization, and accounting. This became a bottleneck for improving throughput on authentication services for large networks.

Benefit: Single Connection is an optimization whereby the NAS maintains a single TCP connection to one or more TACACS+ daemons. The connection is maintained in an open state for as long as possible, instead of being opened and closed each time a session is negotiated. It is expected that Single Connection will yield performance improvements on a suitably constructed daemon.

Considerations: Currently, only the CiscoSecure daemon V1.0.1 supports Single Connection. The NAS must be explicitly configured to support a Single Connection daemon. Configuring Single Connection for a daemon that does not support this feature will generate errors when TACACS+ is used.

Product Marketing Contact: Charles Yager

5.2.2 TACACS+ SENDAUTH Function

Description: SENDAUTH is a TACACS+ protocol change to increase security. SENDAUTH supersedes SENDPASS.

Benefits: The Network Access Server (NAS) can support both SENDAUTH and SENDPASS simultaneously. It detects if the daemon is able to support SENDAUTH, and, if not, will use SENDPASS instead. This negotiation is virtually transparent to the user, with the exception that the down-rev daemon may log the initial SENDAUTH packet as unrecognized.

Considerations: SENDAUTH functionality requires support from the daemon, as well as the NAS.

Product Marketing Contact: Charles Yager

5.3 Kerberos V Client Support

Description: This feature provides full support of Kerberos V client authentication, including credential forwarding.

Benefits: Customers with existing Kerberos V infrastructures can use their Key Distribution Centers (KDCs) to authenticate end-users for network or router access.

Considerations: This is a client implementation, not a Kerberos KDC. Kerberos is generally considered a legacy security service, and is most beneficial in networks already using Kerberos.

Product Marketing Contact: Elizabeth Kaufman

6 Network Management

6.1 HTTP Server

Description: Cisco IOS software introduces an HTML management tool. This tool allows customers to navigate through the command line interface via Web-like hot links. Customers can monitor their routers through an HTML interface to the CLI. Customers can also modify their Web page to add frequently used hot links or to add their company logo.

HTTP Server on a Cisco 7200 series router provides a logical view of the hardware configuration. Customers can point and click on interfaces to check their status or to modify the configuration.

Product Marketing Contact: Bob Berlin

6.2 ClickStart

Description: ClickStart is a powerful new Web-based software solution that enables users to install a Cisco router in minutes. The ClickStart enables Cisco 1000 series ISDN access routers to be accessed by any Web browser on any desktop platform including MS Windows, Windows 95, Windows NT, Unix and MacOS. The easy-to-use Web-based interface guides users through the router installation process. By completing an initial setup form, a user can easily configure the router and bring up the ISDN network connection. The router is then manageable from a central location, so that fine-tuning and upgrades can be performed remotely.

The same interface is used to access the HTTP Server functions, described in Section 6.1.

Product Marketing Contact: Bob Berlin

6.3 MIBs Supported

6.3.1 APPN DLUR MIB

Please see Section 4.1.2 for details.

6.3.2 AtoM MIB Support

Please see Section 3.3.6 for details.

6.3.3 RTTMON Support

Please see Section 4.5 for details.

6.3.4 Cisco IP Encryption MIB

6.3.5 Cisco Modem Management MIB

6.3.6 Cisco SYSLOG MIB

6.3.7 Cisco TN3270 Server MIB

6.3.8 SNA NAU

6.3.9 Cisco Memory Pool

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
World Wide Web URL:
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe s.a.r.l.
Parc Evolic - Batiment L1/L2
16, Avenue du Quebec
BP 706 - Villebon
91961 Courtaboeuf Cedex
France
Tel: 33 1 6918 61 00
Fax: 33 1 6928 83 26

Intercontinental and Latin American Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
Tel: 408 526-7660
Fax: 408 526-4646

Japanese Headquarters
Nihon Cisco Systems K.K.
Seito Kaikan 4F
5, Sanbancho, Chiyoda-ku
Tokyo 102
Japan
Tel: 81 3 5211 2800
Fax: 81 3 5211 2810

Europe, Middle East, and Africa

Austria
Cisco Systems Austria GmbH
World Trade Center
A-1300 Vienna Airport
Austria
Tel: 43 1 7007 6256
Fax: 43 1 7007 6027

Belgium
Cisco Systems Brussels
Complex Antares
71 Avenue des Pleiades
1200 Brussels
Belgium
Tel: 32 2 778 42 00
Fax: 32 2 778 43 00

Denmark
Cisco Systems
Larsbjoernsstraede 3
DK-1454 Copenhagen K
Denmark
Tel: 45 33 37 71 57
Fax: 45 33 37 71 53

Finland
Cisco Systems
Maistraatinportti 2A
FIN-00240 Helsinki
Finland
Tel: 358 1594 3090
Fax: 358 1594 3093

Germany
Cisco Systems GmbH
Max-Planck-Strasse 7, 3rd Floor
85716 Unterschleissheim
Germany
Tel: 49 89 32 15070
Fax: 49 89 32 150710

Ireland
Cisco Systems Ltd.
Europa House, 4th Floor
Harcourt Street
Dublin 2
Ireland
Tel: 35 3 1 475 4244
Fax: 35 3 1 475 4778

Italy
Cisco Systems Italy Srl
Centro Direzionale Milano Oltre
Palazzo Raffaello Scala B 4P
Via Cassanese 224
20090 Segrate (Mi)
Italy
Tel: 39 2 26 97 31
Fax: 39 2 26 92 9006

The Netherlands
Cisco Systems
Stephensonweg 6
4207 HB Gorinchem
The Netherlands
Tel: 31 183 622 988
Fax: 31 183 622 404

Norway
Cisco Systems
Holmens Gate 4
N-0250 Oslo
Norway
Tel: 47 22 83 06 31
Fax: 47 22 83 22 12

Portugal
Cisco Systems Portugal
Avda. da Liberdade 114-134
1250 Lisboa
Portugal
Tel: 351 1 340 45 31
Fax: 351 1 340 45 75

South Africa
Cisco Systems South Africa
Meintjie Parker House
328 Rivonia Blvd.
Rivonia, Gauteng
South Africa
Tel: 27 11 807 4444
Fax: 27 11 807 4447

Spain
Cisco Systems Spain
Avenida de Burgos, 17 Pl. 11
Edificio Triada II
28036 Madrid
Spain
Tel: 34 1 383 2178
Fax: 34 1 383 8008

Sweden
Cisco Systems AB
Arstaangsvagen 13
S-117 60 Stockholm
Sweden
Tel: 46 8 681 41 60
Fax: 46 8 19 04 24

Switzerland
Cisco Systems Switzerland
Grossrietstrasse 7
CH-8606 Naenikon/ZH
Switzerland
Tel: 41 1 905 20 50
Fax: 41 1 941 50 60

United Arab Emirates
Cisco Systems (Middle East)
PO Box 26095
City Tower 2
Sheik Zayed Road
Dubai, UAE
Tel: 971 4 318 788
Fax: 971 4 313 681

United Kingdom
Cisco Systems Ltd.
4 New Square
Bedfont Lakes
Feltham, Middlesex TW14 8HA
UK
Tel: 44 1 81 818 1400
Fax: 44 1 81 893 2824

Intercontinental

Argentina
Cisco Systems Argentina
Cerrito 1054, Piso 9
(1010) Buenos Aires
Argentina
Tel: 54 1 811 7526
Fax: 54 1 811 7495

Australia
Cisco Systems Australia Pty Ltd
Level 17
99 Walker Street
North Sydney NSW 2060
Australia
Tel: 61 2 9935 4100
Fax: 61 2 9957 4077

Brazil
Cisco Systems Do Brasil
Rua Helena 218, 10th Floor
Cj 1004-1005 Vila Olimpia
Sao Paulo, SP CEP 04552-050
Brazil
Tel/Fax: 55 11 822 6095
Tel/Fax: 55 11 822 6396

Canada
Cisco Systems Canada Limited
150 King Street West
Suite 1707
Toronto, Ontario M5H 1J9
Canada
Tel: 416 217-8000
Fax: 416 217-8099

Central America / Caribbean
Cisco Systems, Inc.
790 NW 107th Avenue, Suite 102
Miami, Florida 33172
USA
Tel: 305 228-1200
Fax: 305 222-8456

Chile
Cisco Systems-Chile
Avenida Tajamar 481, Ofi. #101
Las Condes
Santiago, Chile
Tel: 562-339-7000
Fax: 562-339-7022

China, PRC
Cisco Beijing Representative Office
Add. Unit 751
New Century Hotel Office Tower
No. 6 Southern Road Capital Gym
Beijing, 100044
China, PRC
Tel: 86-10-8492398
Fax: 86-10-8492395

Colombia
Cisco Systems Colombia
Cra. 18 #86A-14
Bogota
Colombia
Tel: 57 1 296 0067
Fax: 57 1 616 3030

Costa Rica
Cisco Systems-Costa Rica
De las Tunas 100 metros Norte
50 metros Este
Sabana Norte
San Jose
Costa Rica
Tel: 506-296-1885
Fax: 506-296-3607

Hong Kong
Cisco Systems (HK) Ltd
Suite 1009, Great Eagle Centre
23 Harbour Road
Wanchai
Hong Kong
Tel: 852 2583 9110
Fax: 852 2824 9528

India
Cisco Systems (HK) Ltd
New Delhi Liaison Office
Suite 119, Hyatt Regency Delhi
Bhikaiji Cama Place, Ring Road
New Delhi 110 066
India
Tel: 91-11-616-7688
Fax: 91-11-616-7688

Indonesia
Cisco Systems, (HK) Ltd
Level 12, Wisma Bank Dharmala, JI
Jenderal Sudirman Kav. 28
Jakarta Selatan 12910
Indonesia
Tel: 62 21 523 9132
Fax: 62 21 523 9259

Korea
Cisco Systems Korea
Samik Rabidol Building 5th floor
720-2 Yuksam-2-dong, Gangnam-ku
Seoul, 135-082
Korea
Tel: 82 2 3453 0850
Fax: 82 2 3453 0851

Malaysia
Cisco Systems (Malaysia) Sdn. Bhd.
14.05, 14th Floor
Menara Multi-Purpose, Captial Square
8 Jalan Munshi Abdullah
50100 Kuala Lumpur
Malaysia
Tel: 0203-292-8398
Fax: 0203-292-8389

Mexico
Cisco Systems de México, S.A. de C.V.
Ave. Ejecito Nacional No. 926
3er Piso
Col. Polanco C.P. 11560
Mexico D.F.
Tel: 52 5 328 7600
Fax: 52 5 328 7699

New Zealand
Cisco Systems New Zealand
Level 16, ASB Bank Centre
135 Albert Street
PO. Box 6624
Auckland
New Zealand
Tel: 64 9 358 3776
Fax: 64 9 358 4442

Philippines
Cisco Systems Manila Office
The Executive Tower Centre
Room 9, 24/F, Pacific Star Building
Sen. Gil J. Puyat Cornet
Makati Avenue, Makati City
Metro Manila
Philippines
Tel: 632 892 4476
Fax: 632 811 5998

Singapore
Cisco Systems (USA) Pte Ltd
501 Orchard Road
#04-11 Lane Crawford Place
Singapore 238880
Tel: 65 738 5535
Fax: 65 738 2202

Taiwan, ROC
Cisco Systems (HK) Ltd
Taiwan Representative Office
16F/B 333 Tunhua South Road
Sec 2
Taipei
Taiwan, ROC
Tel: 886-2-738-6667
Fax: 886-2-377-0611

Thailand
Cisco Systems (HK) Ltd
7th Floor, The Park Place Building
231 Sarasin Road, Pathumwan
Bangkok 10330
Thailand
Tel: 662 253 5315
Fax: 662 253 8440

Venezuela
Cisco Venezuela
Calle Bajada de Los Curtidores
Qta. Jakaranda - Alto Hatillo
Caracas
Venezuela
Tel/Fax: 58 2 963 6140

United States

Central Operations
5800 Lombardo Center
Suite 160
Cleveland, OH 44131
Tel: 216 520-1720
Fax: 216 328-2102

Eastern Operations
1160 West Swedesford Road
Suite 100
Berwyn, PA 19312
Tel: 610 695-6000
Fax: 610 695-6006

Federal Operations
380 Herndon Parkway
Herndon, VA 22070
Tel: 703 397-5500
Fax: 703 397-5599

Northeastern Operations
One Penn Plaza
Suite 3501
New York, NY 10119
Tel: 212 330-8500
Fax: 212 330-8505

Northern Operations
8009 34th Avenue South
Suite 1550
Bloomington, MN 55425
Tel: 612 851-8300
Fax: 612 851-8311

Service Provider Operations (Telecommunications)
111 Deerwood Road
Suite 200
San Ramon, CA 94583
Tel: 510 855-4800
Fax: 510 855-4896

Southwestern Operations
14160 Dallas Parkway
Suite 400
Dallas, TX 75240
Tel: 214 774-3300
Fax: 214 774-3344

Western Operations
2755 Campus Drive
Suite 205
San Mateo, CA 94403
Tel: 415 377-5600
Fax: 415 377-5699

Cisco Systems has more than 125 sales offices worldwide. To contact your local account representative, call Cisco's corporate headquarters (California, USA) at 408 526-4000 or in North America, call 800 553-NETS (6387).

0496R

(0796R)