WHITE PAPER

CARRIER-CLASS TELEPHONY OVER IP

Ericsson Engine Integral and Cisco IP Backbone

This white paper, developed by Cisco Systems® and Ericsson, describes the requirements that telephony over IP impose on IP networks and how to fulfill these requirements. The paper focuses on Ericsson's ENGINE Integral Softswitch Solution applied to a core network based on Cisco® 12000 Series routers. This paper is coauthored by market leaders working together to provide best-in-class solutions to deliver carrier-scale, carrier-class, softswitch solution infrastructure.

INTRODUCTION

Transport of carrier-class telephony over IP places stringent requirements on the design of an IP network. In particular, low delay, low jitter, low packet loss, protection from denial-of-service attacks, and high availability are of paramount importance. The design characteristics and the architectural choices that fulfill these requirements may differ in various service provider environments, all depending on the scaling requirements, physical and routing topology of the existing network, relation to traditional services, manageability, etc. There are numerous IP network design options available to service providers. The design concepts described in this document are tested and verified in Ericsson's proof-of-concept lab in collaboration with Cisco Systems. References to configuration guides and verification results are listed at the end of this document.

The scope of this paper includes the following:

- Terminology
- Functional overview of Ericsson ENGINE Integral solution
- Functional overview Cisco 12000 Series Router
- Telephony requirements
- Solution A: MPLS traffic engineering in respect to resilience, class of service (CoS), and network security discussion
- Solution B: Fast IP design in respect to resilience, CoS, and network security
- Solutions summary discussion
- Common Cisco IOS® Software tools
- References and further reading

This paper does not discuss Internet service provider (ISP) best practices or protocol design in depth and the reader is expected to be familiar with those concepts.

This paper is intended for networking professionals. To fully understand and apply the information in this paper, the reader should already be familiar with Cisco IOS Software, the Cisco 12000 Series architecture, The ENGINE Integral architecture, ISP essentials, and routing design concepts.

BACKGROUND

Ericsson's ENGINE Integral solution has initially been deployed with SONET/SDH and ATM infrastructures as transport between media gateways and signaling servers, often referred to as softswitches. However, the ENGINE Integral architecture is created to allow pure packet transport over IP using popular media types including packet over SONET (POS) and Ethernet. This means that existing installations can be upgraded to support IP packet-transport technology. Using open protocols such as H.248 and Session Initiation Protocol for Telephony (SIP-T) allows this solution to easily integrate with and complement other solutions.

Cisco is a leading networking vendor for routing and switching solutions for both service providers and enterprises. Cisco has deployed enterprise IP voice solutions as well as service provider voice solutions in IP and Multiprotocol Label Switching (MPLS) networks based upon carrier-class technology.

Customers that are overlapping their installed base will directly benefit from the combined knowledge of Ericsson and Cisco, but so too will customers planning new deployments or upgrading other existing installations.

IP networks have long been synonymous with the Internet, but there are many mission-specific networks based on IP deployed today. Carrier-class routers have facilitated the integration of these special-purpose networks into common multiservice-enabled core networks. This has been done with the use of VPN technology, enhanced security features, and strict CoS implementations. Many network operators have selected to use MPLS, but MPLS is not a mandated requirement.

This paper is intended to assist customers making deployment selections today. Every network is unique and to best match each customer's requirements, a customer-specific deployment plan should be discussed with specialists from Ericsson and Cisco.

TERMINOLOGY

The following terms are explained briefly to help avoid confusion.

- *Media gateway*—The main task of the media gateway is to provide media interworking between the circuit-switched domain and the IP domain. They forward traditional telephony signaling toward the telephony servers and convert voice samples to IP packets. The media gateway is controlled by the telephony server.
- *Telephony servers*—These are media gateway controllers (MGCs). They control the media gateways to provide traditional telephony services. Each telephony server controls one or more media gateway, but a single media gateway cannot be controlled by multiple telephony servers. In cases where multiple telephony servers exist, they communicate via SIP-T or Bearer Independent Call Control (BICC) signaling protocol.
- *End systems*—In this context, softswitch nodes comprise media gateways, telephony servers, and other telephony-related devices, providing carrier-grade telephony services such as public switched telephone network (PSTN) and IP-based telephony such as H.323 and SIP.
- *Telephony aware*—This term is used to indicate if an IP router is carrying the IP addresses that represent the telephony end systems.
- *Domain*—A subset of a softswitch network created to partition the larger network into smaller components. One telephony server is expected per domain and one domain can span several IP areas and autonomous systems.

- *Telephony edge router*—This is an ordinary IP router normally located at the edge of the IP backbone network and in front of end systems. Telephony edge routers employ Border Gateway Protocol (BGP) policy and access filtering rules to protect end systems from DoS attacks and intrusion. The telephony edge router defines the softswitch perimeter and handoff.
- *Telephony area border (TAB) routers*—These routers are area-transit routers that create a hierarchical overlay to connect telephony domains across multiple autonomous systems. They can be used for both MPLS and IP designs. To bind together domains, they have to be telephony aware.
- *Transit router*—These are core-network transit routers and they are used to interconnect telephony edge routers. Transit routers are not telephony-aware.

ERICSSON ENGINE INTEGRAL—FUNCTIONAL OVERVIEW

Figure 1 depicts the network overview and the components used to support telephony over IP (ToIP) with ENGINE Integral. ISUP, DSS1, and V5.2 are protocols used for signalling to Ericsson's Remote Subscriber Stage (RSS) and EngineAccessRamp (EAR). (AN: access node, LE: local exchange, TeS: telephony server.)

Figure 1
Overview of ENGINE Integral Network

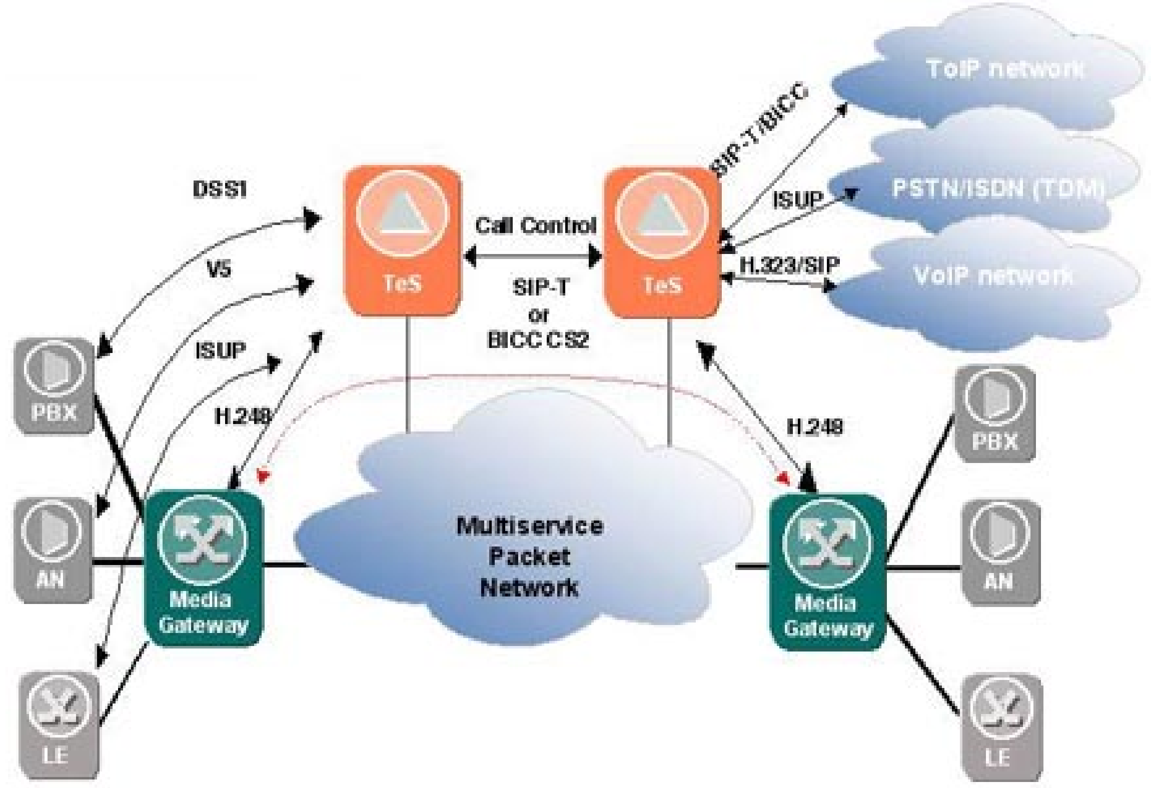

The main building blocks of ENGINE Integral are the telephony server and the media gateway, which communicate via H.248 protocol. The telephony server handles call logic and controls the switching resources that are implemented in the media gateway. The media gateway performs the real switching and media interworking functions between the circuit switching and IP domains. Local exchanges, transit exchanges, access nodes, the ENGINE Access Ramp, and ISDN PBXs may be connected to the media gateways.

This means that the telephony server operates on the transit level as well as on the local level. Each telephony server controls one or more media gateways. All the media gateways controlled by one telephony server define a server domain. However, a given media gateway cannot be controlled by more than one telephony server. In case of multiple telephony servers, communication between these takes place via SIP-T or BICC signaling. A media gateway generates the following flows: signaling, OAM, and voice (Figure 2). This solution allows for adoption of other gateway types and signaling protocols, but that is not covered in this document. For the purpose of filtering and specific forwarding operation, each traffic flow from the media gateway will be assigned a different address range.

Figure 2
Ericsson's Media Gateway IP Structure and Flow Components

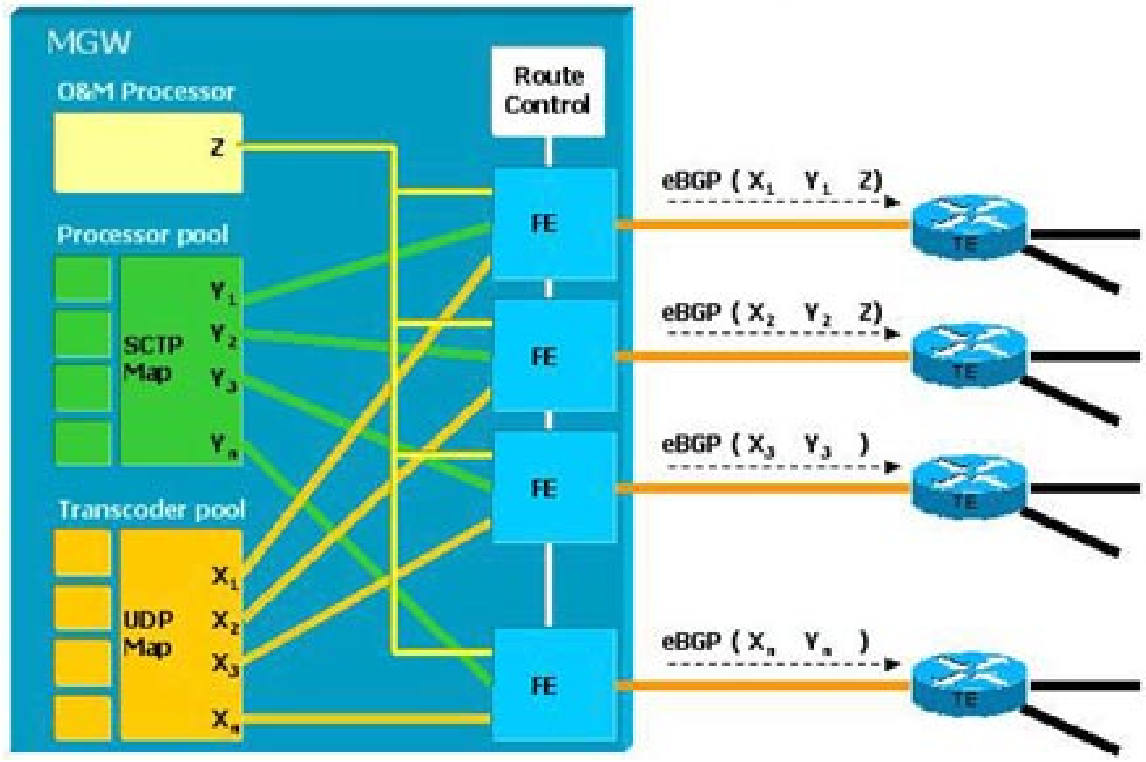

Signaling (Media Gateway Control Protocol)

Media Gateway Control Protocol (MGCP), also known as H.248, is carried over Stream Control Transmission Protocol, SCTP. SCTP transport address is defined as combination of an IP address and SCTP port number. SCTP supports multihoming capability for redundancy, and as a result SCTP has two IP addresses. In the event of link or node failures, SCTP will use the secondary link until the primary path is available again. The SCTP number identifies the SCTP termination point (for the receiving process). The media gateway announces its address pairs on different interfaces. Thus the two signaling paths should have maximum diversity through the network. Most of the control messages have around 200 bytes, and only a minor proportion of packets can be up to 1500 bytes. The SCTP packets are loss-sensitive but not very delay-sensitive.

Operation, Administration, and Maintenance

Operation, administration, and maintenance (OAM) traffic refers to Secure Shell (SSH) Protocol, Simple Network Management Protocol (SNMP), FTP, Network Time Protocol (NTP), and HTTPS type of traffic for the purpose of configuration and supervision of different network elements. The OAM traffic has its own IP address, which is announced via two media gateway interfaces. This address can be private or public. The OAM traffic is not as sensitive to delays as packet loss.

Voice and Backhauled Signaling

This type of traffic is carried via Real-Time Transport Protocol/User Datagram Protocol (UDP) encapsulation. The RTP/UDP sessions are uniquely identified via the combination of an IP address and UDP port number. Voice payload IP packets use one of the two virtual addresses (X1 and X2 in Figure 2). These addresses are then tied to each IP interface address. For instance, the X1 and Y1 addresses are only accessible from IP interface 1, and similarly X2 and Y2 addresses are only accessible from IP interface 2. The UDP port number identifies the correct codec, which handles the voice channel. A voice packet consists of 160-byte voice sampling (assuming G.711 encoding) plus a 40-byte RTP, UDP header. This type of traffic is highly sensitive to delay and jitter but not so sensitive to loss. To reduce the impact of interface or link failure on the media gateway, the load will be distributed on a number of interfaces. In case one of the interfaces fails, only a portion of the traffic will be affected. New incoming calls will be forwarded only to the working interface and link. This procedure does not provide redundancy but it does minimize the effects of a failure.

CISCO 12000 SERIES ROUTER—FUNCTIONAL OVERVIEW

The Cisco 12000 Series Router delivers capacity and services with its fully distributed forwarding architecture and maximum-efficiency crossbar switch fabric. The combination of the centralized scheduler and unique virtual-output queuing technology maximizes the use of the switch-fabric bandwidth, minimizes latency, and provides nonblocking performance. The Cisco 12816 Router uses the latest in high-performance, application-specific integrated circuit (ASIC) technology to provide line-rate forwarding with an extensive feature set, while maintaining the strict jitter and latency required for real-time services. Offering a comprehensive set of quality of service (QoS), Multiprotocol Label Switching (MPLS), and high-availability features, the Cisco 12000 Series can help ensure maximum bandwidth utilization and traffic differentiation while meeting even the strictest customer service-level agreements (SLAs).

Figure 3 shows the Cisco 12000 Series Router Architecture, displaying the split between the control and data planes.

Figure 3
Cisco 12000 Series Router Architecture Overview

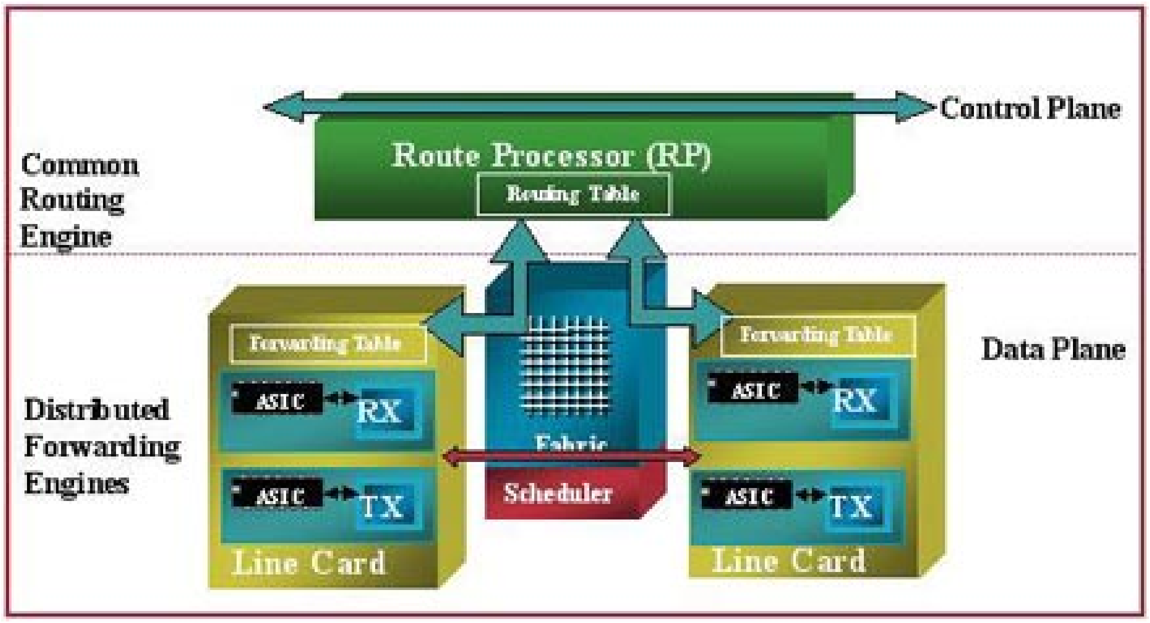

This is complemented with a large set of Cisco IOS Software functions that optimize routing protocol behavior and resilience.

Tight SLA designs are made possible with the layered structure of queues, shapers, and policers that the Cisco 12000 Series has on all contention points. Traffic can be classified, rate limited, and shaped on ingress, then queued according to classification toward the fabric. Low-Latency Queuing (LLQ) helps ensure that priority traffic is scheduled first. Each egress line card can then further meter and classify traffic and perform class-based queuing on shaped bundles. Each line card has up to 4000 queues for ingress processing and an equal amount for egress port or class assignment. This helps ensure priority delivery across a network of routers.

Figure 4 shows how gray packets are received by an ingress port and then classified and rate controlled per service class, before being scheduled to the destination slot or port. The egress line card can further classify and process packets before they are transmitted out of an interface.

Figure 4
Packet Classification and Processing

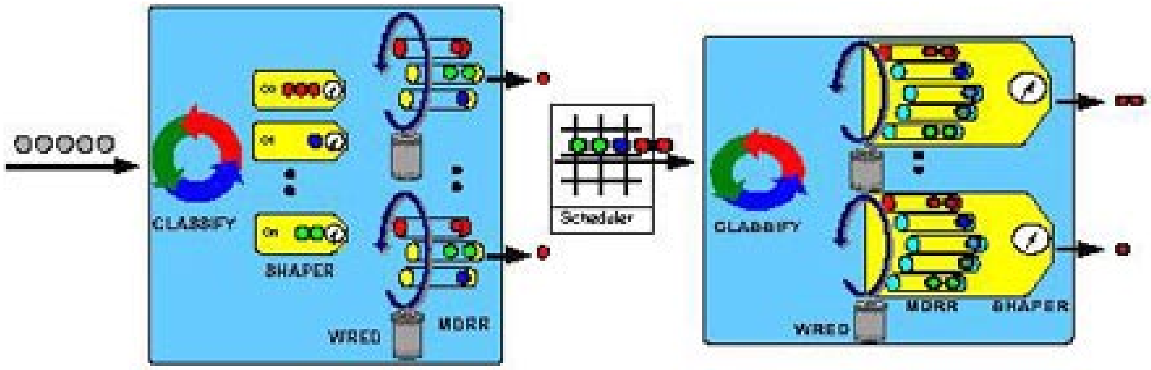

TELEPHONY REQUIREMENTS

As described in the ENGINE Integral section, packet voice is sensitive to delay, jitter, and loss. The requirements on running voice over IP (VoIP) relates to two things; the users experience and the applied protocol's robustness. The ENGINE Integral components can adapt to loss and delay beyond the user experience recommendations:

The voice user's experience directly relates to the following requirements:

- Low packet loss for voice (less than one percent).
- Low delay; mouth-to-ear delay should ideally be no more than 150 milliseconds (ms), including codec, access links, and core network. If delays can be kept below 150 ms, most applications, both speech and nonspeech, will experience essentially transparent interactivity (G.114). The latency components are described in more detail in Appendix C.
- Average jitter should be in the range of 30–50ms. Jitter less than 50 ms is possible with core links above E1 or T1 rates (Y.1541).

Although packet loss of any kind is undesirable, some loss can be tolerated. Some amount of packet loss for voice services could be acceptable as long as the loss is spread out over a large amount of users. A well-known measure used to determine the voice quality is called Mean Opinion Score (MOS). Figure 5 illustrates the relationship between packet-loss ratio and the MOS. In essence this means that the service quality will dictate codec selection and set the network's packet-loss requirement.

Figure 5
Mean Opinion Score as a Function of Packet Loss

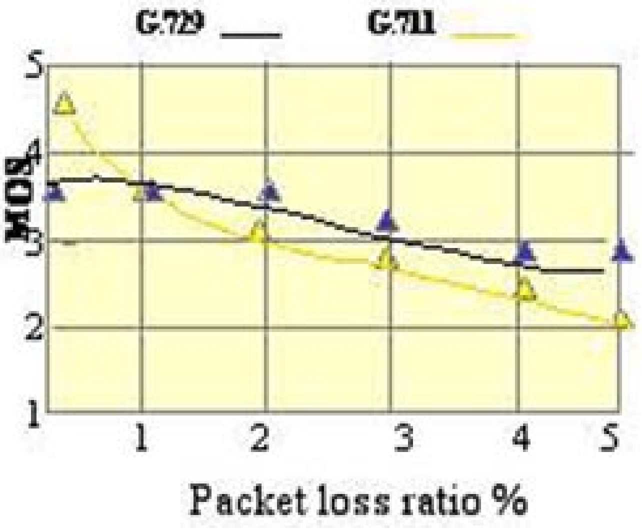

A packet-loss concealment technique can tolerate packet loss up to 0.5 percent (G.711). Loss of voice packets may cause gaps known as “temporal clipping.” According to ITU-T G.116 the temporal clipping segments (for example, the time clipping prevails) should be under 64 ms and below 0.2 percent of the active speech.

Requirements for the IP backbone network can be summarized as follows:

- The IP network should recover from link or node failures in less than 2 seconds to avoid user distress or expiration of signaling protocol timers. This will imply a protection scheme that is slightly better than standard IP.
- Packet loss should be less than 1 percent. This may or may not require a strict SLA, depending on current traffic patterns.
- Core network latency below 10 ms is achievable if each node introduces less than 0.5 ms in latency and jitter. Transmission delay is not included.
- The network should include a media gateway that announces IP addresses from a particular interface, with voice traffic arriving on the same interface.
- Telephony end systems must be protected from DoS attacks and IP spoofing.
- The network should always be highly available.

SOLUTION A: MPLS TRAFFIC ENGINEERING

Figure 6 shows the active components of a voice domain: media gateways, telephony servers (TeS), and telephony edge routers (TE). Each end system is connected to a pair of telephony edge routers, which can support multiple end systems.

Figure 6
One Voice Domain with Signaling Servers and Media Gateways

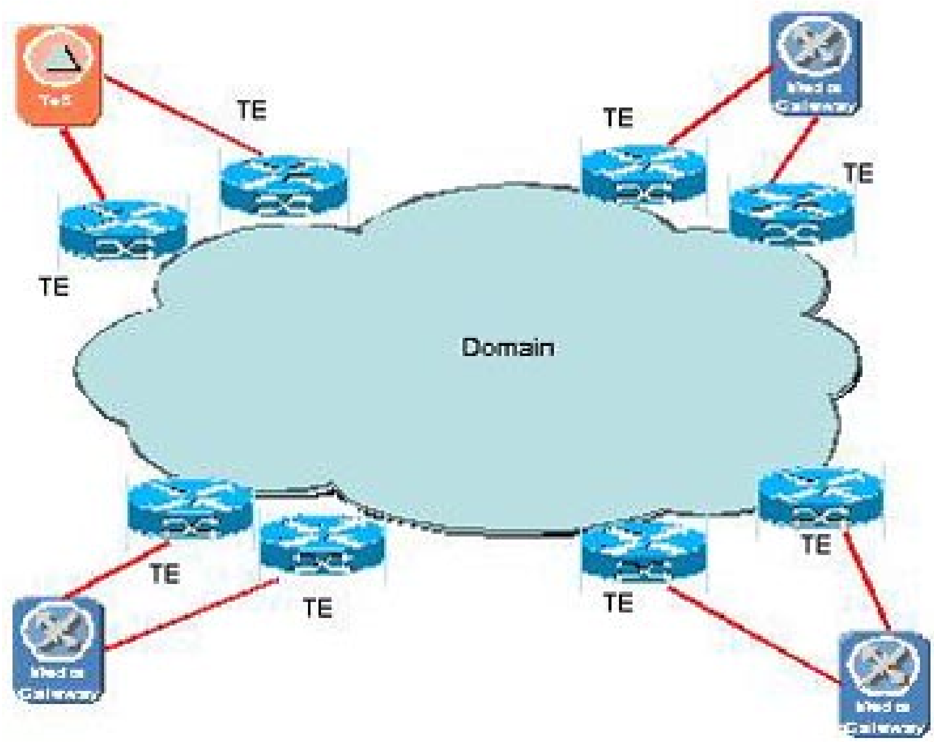

The tunnel infrastructure has the following characteristics:

- Every telephony edge router has one tunnel defined to both routers in every remote pair. This helps ensure that every telephony edge router can reach every other telephony edge router and keeps the telephony end system's address space hidden from the core network.
- All telephony edge routers use internal BGP (iBGP) to peer.
- Each tunnel has two path options: primary and secondary. Because only one path is preestablished, the number of active tunnels is conserved. While failing over from one tunnel interface to another requires IGP to converge, path errors require RSVP to detect the failure and establish a secondary backup path.
- In total, there are 4N(N-1) tunnels per domain, excluding interdomain tunnels. Thus 40 end systems would result in less than 1500 tunnels and active paths.

Figure 7 shows primary and secondary path options displayed between two telephony edge routers (TE1 and TE3).

Figure 7
Primary and Secondary Path Options

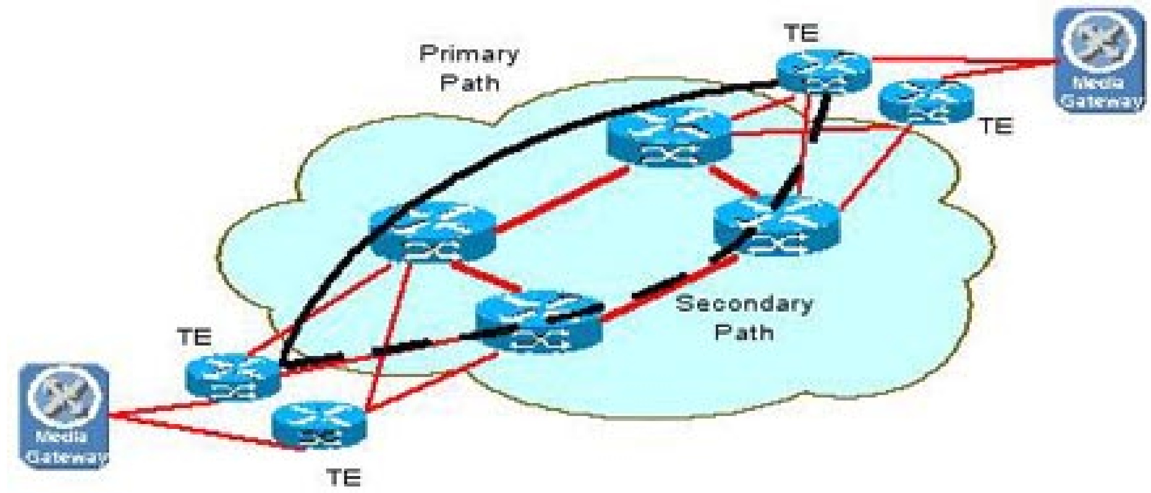

A slightly different technique is required to scale beyond one domain and to span across autonomous system boundaries and by building an overlay of Telephony Area Border routers, the network can scale for very large installations. In addition to the intradomain mesh of tunnels, additional tunnels are needed between each telephony edge router and the TAB routers.

Figure 8 shows an interdomain tunnel design, including one exit tunnel with two path options per tunnel edge router.

Figure 8
Interdomain Tunnel Design

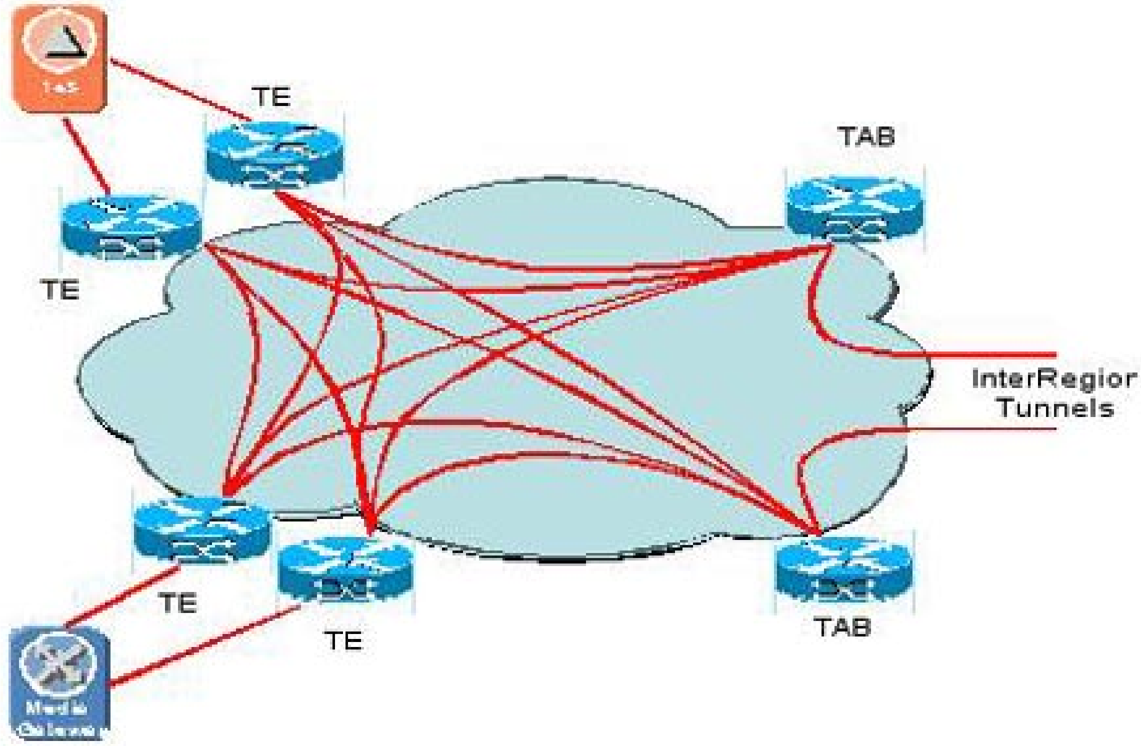

To interconnect domains, the TAB routers are connected with a mesh of tunnels.

Figure 9 shows how TAB routers are used to interconnect domains. Only the tunnels of domain 1 are shown.

Figure 9
Interconnected Domains

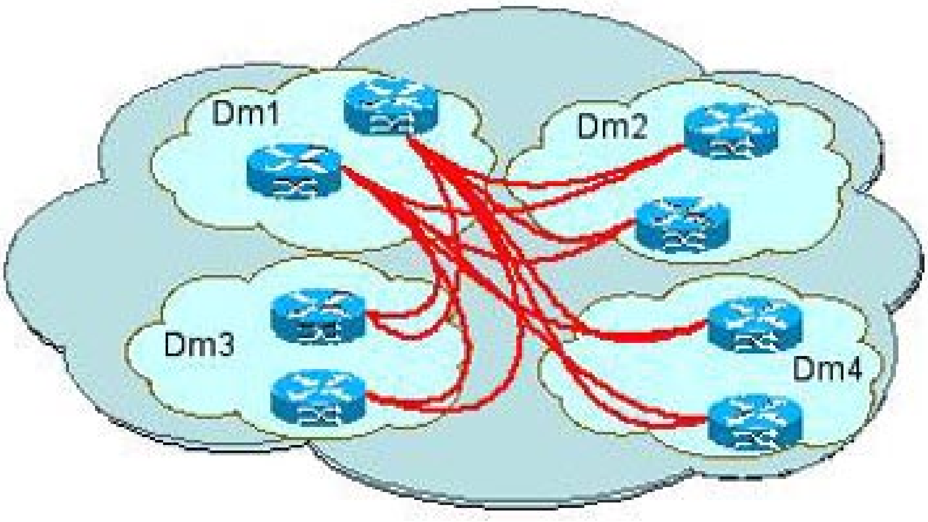

Convergence

This solution depends on RSVP path errors for convergence and this should be well within 2 seconds. Increasing the number of tunnels and high background network "stress" has a small impact, but subsecond convergence should still be expected. MPLS Fast Reroute further improves failover times, but is not needed to fulfill this requirement.

Network Security

Because the core network nodes are not telephony-aware, they cannot send packets to the media gateways. The telephony edge routers are telephony-aware and need to protect end systems with access lists and Unicast Reverse Path Forwarding (uRPF) checking. As always, all intermediate routers need to be protected from intrusion and protocol-specific spoofing. More details on security aspects are described in Appendix B.

Class of Service

Latency and jitter requirements can be satisfied with the use of Modified Deficit Round Robin (MDRR) queuing and depending on other core service classes, LLQ may be applied. End systems control packet classification and the core network is expected to ensure priority delivery of voice packets. Further, the core network must implement classification and access control for nontrusted sources.

SOLUTION B: NATIVE FAST IP

The IP design differs from the MPLS design in the following ways:

- It depends on the IGP for restoration
- It does not require any component of MPLS to run
- The core routers need to be telephony aware
- There is no scale impact by the number of end nodes per domain

The active components of the IP solution are the same as for solution A and are previously displayed in Figure 6.

Good IGP convergence is directly related to design and scale. Designs that allows for load sharing through parallel links or other equal-cost paths make a link state routing protocol converge faster. In addition, the following timers and behavior can be tuned:

- Carrier detection
- POS alarm processing
- Partial support of Shortest Path First (SPF)
- Incremental SPF support
- Flooding behavior
- Link-state packet (LSP) flooding timers
- SPF timers
- Routing table convergence priority

Controlling these components will help enable an IGP such as Intermediate System-to-Intermediate System (IS-IS) Protocol or Open Shortest Path First (OSPF) Protocol to convergence in sub-2-second times. As networks grow and more nodes and addresses are added, it is essential to ensure that the IP addresses used by the telephony end systems are tagged for priority convergence. This function is available in Cisco IOS Software Release 12.0(26)S and later. In short, it allows the routing process to recognize a subset of addresses as more important than other addresses and process these first. This results in predictable convergence even in growing networks.

The principal routing design of a native fast IP network includes the following characteristics:

- Media gateways use external BGP (eBGP) to local telephony edge routers
- Telephony edge routers share IGP with IP core network and redistribute local IP telephony routes
- Telephony edge routers use iBGP toward TAB routers for central “telephony-peering” policy control
- Optionally, TAB routers run eBGP toward autonomous system boundary routers (ASBRs) for simplified handoff; otherwise the ASBRs control the policy instead of the TABs
- IP core network uses only loopback and links local addresses in IGP and iBGP for all other routes

Figure 10 shows a routing design that differs from Solution A only in that there is no direct iBGP peering between the telephony edge routers. The ASBRs or optional TAB routers will act as route reflectors for the telephony edge routers and control the policy interaction between the IP core network and the telephony network.

Figure 10
IP Network Design Overview

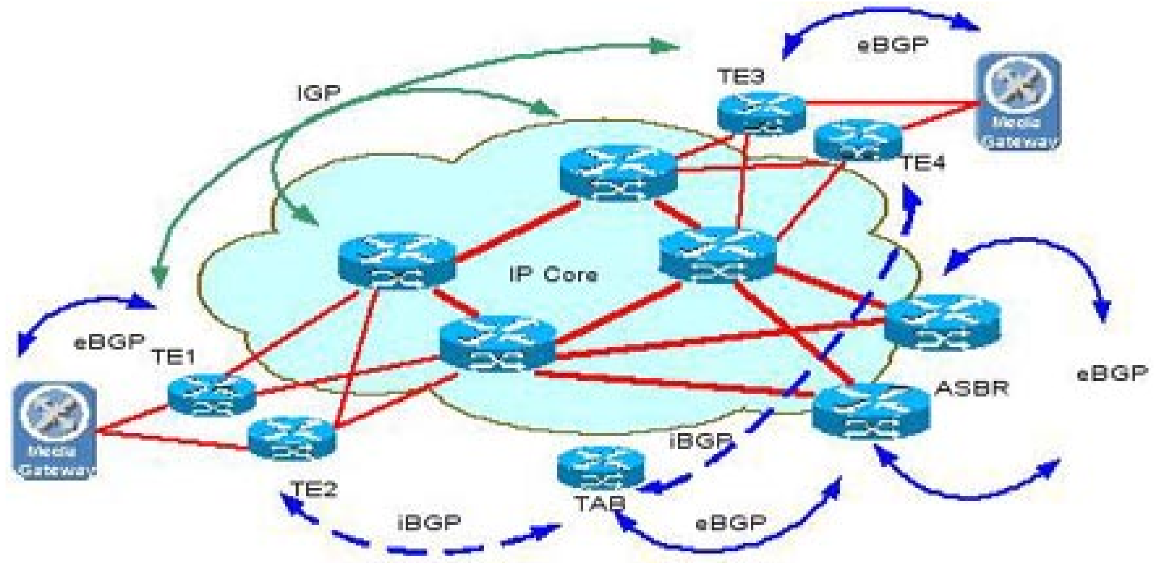

Convergence

Optimized IGP configuration is essential for predictable behavior, but the actual topology has impact and careful design may vastly improve network behavior. An IGP is expected to only carry router identification or loopback addresses, and link addresses. This allows even very large networks to have small IGP tables. This solution requires the telephony end-system addresses to be carried within the IGP, and with few addresses per media gateway, this should not be an issue. IGP parameter tuning is needed to fulfill the convergence requirement.

Network Security

All routers in a domain need to be telephony-aware, hence the telephony edge routers must actively protect the telephony end systems. The protocol design will help ensure that the telephony edge routers only have telephony and core IP addresses in their tables, similar to Solution A. The security aspects are therefore similar to those of the MPLS telephony edge solution.

Class of Service

Class of service is the same as in Solution A: Latency and jitter requirements can be satisfied with the use of MDRR queuing and, depending on other core service classes, LLQ may be applied. End systems control packet classification and the core network is expected to ensure priority delivery of voice packets. Further, classification and access control for nontrusted sources must be implemented in the core network.

SOLUTIONS SUMMARY

Ericsson's ENGINE Integral Softswitch solution applied to a Cisco 12000 based core can support network design based up on either IP or MPLS while delivering deterministic service behavior. It is important that both solutions can be supported, since customers' business models and strategies often are reflected in their selection of network architecture.

The primary requirements of the IP telephony network design are:

- High availability
- Quality of service
- Network security

As mentioned earlier, although there are a number of architectural options to satisfy these requirements, this document has focused on two.

Solution A relies on RSVP-based traffic engineering tunnels in combination with differentiated services. MPLS offers a method to establish connection-oriented paths or LSPs in an otherwise connectionless IP network. The major argument for using RSVP is that a native IP network is more difficult to engineer to enable path diversity and to provide precise control over how traffic is rerouted when the primary path is faced with single or multiple failures. It should be noted that the added element of control that is achieved with MPLS traffic engineering will come with an additional cost in the form of complexity. The overall complexity is relative to the experience that the service provider has. Some may already have extensive MPLS deployment for other services such as Layer 3 VPNs, Layer 2 transport, Layer 2 VPNs using RSVP, or LDP. The mechanisms described in this document do not exclude those services to be deployed in the same backbone network at the same time.

Solution B shows how a tuned IP network can achieve the same requirements. QoS and network security are similar for both solutions. The main benefit of using native IP is simplicity; if it is easy to design and build, it is also easier to manage and troubleshoot.

While the choice to deploy a MPLS or IP-based network may have fundamental impact on network operations and service deployment, the designs discussed in this paper showed that the MPLS design is a superset of the IP design: When building a capable MPLS network, one must first build a capable IP network. There are no shortcuts. The network must have a solid routing design, a well-planned DoS protection scheme, and a QoS design that manages untrusted domains. Once the foundation is laid, one can start looking at service-specific design considerations.

Multiservice networks require cross-service classification alignment; the backbone network must be able to ensure priority traffic at all times. Adding telephony to a network that already carries priority video services will require class consolidation and possibly more bandwidth. On the other extreme, adding enough bandwidth may avoid QoS deployment.

It is feasible to see benefits in allowing the telephony edge router to perform other tasks in parallel, for both the described alternatives. Before doing that, one should consider the following:

- Will the added complexity and additional operating expense compensate for the savings in hardware?
- Do the combined services share the same maintenance windows?
- Are the different services maintained by the same operational staff?
- Are the SLAs of equal magnitude?

Because core network nodes have much less operational changes than edge routers, core sharing is preferred—if any.

Another important aspect of networks today is their ability to protect themselves from DoS attacks and intrusion. Considering the telephony service described in this document, it is important to ensure that the telephony end systems are treated like a closed user group. The tools that allow for this group to be created are listed in Appendix A. For designs deviating from the recommendations in this document, further tools should be considered.

The path from a best-effort IP network and TDM based voice services to multi-service capable core using softswitches, can take many shapes. But Carrier Class Telephony IP networks are being built today and with the support from vendors such as Cisco Systems and Ericsson, professional program management is available to enable your network.

APPENDIX A. COMMON CISCO IOS SOFTWARE TOOLS

This section lists tools that are available in IP routers. Some of them are likely to be used in the preceding designs, others are optional. They could all be applied differently to allow for other designs.

BGP Communities

Routing protocols can tag routes for local or remote processing. This is useful when trying to separate routes with different origins and purposes; for example, to differentiate between Internet routes and telephony-aware routes.

The MPLS telephony edge solution described in Solution A does not require the use of communities—unless the telephony edge routers are selected to handle other nontelephony traffic. In such scenarios, tags can be used for features like BGP Policy Propagation to limit traffic from non-PSTN sources to zero packets per second.

BGP Shortcuts

While not being a specific feature, shortcuts are used to denote a technique where BGP is used over an abstraction layer. The abstraction has the benefit that none of the intermediate routers need to be telephony-aware. MPLS does this with the use of RSVP or traffic engineering tunnels, which is used in Solution A. Tunneling techniques are available for IP as well, but are not used in Solution B.

BGP Route Filtering and Policy

While BGP Communities can be used to tag routes, the function that applies routing policy to Cisco protocols is yet another function. A policy is applied through a sequence of matching and set commands, where attributes are set in correspondence with the policy by matching criteria. This is first applied on all external peering links for most IP networks, but this can be further detailed to include voice-peering policy. Peering-policy recommendations are beyond the scope of this document, which assumes that all voice peering is done between domains under the same administration.

IGP Routes and Filtering

Link-state protocols assume that all nodes in an area have the exact same view of the network, hence route filtering should not be done on a node basis. To maximize convergence performance, the IGP's database should be small. Experience has shown that one should avoid having much more than the local-link addresses and router-loopback interfaces in the IGP. Because the loopbacks are used as BGP nexthops, this is enough for full reachability without performance impact.

Reasonable IGPs are IS-IS and OSPF.

Solution B relies on the IGP route-tagging feature to identify the networks that should be treated with high priority when processing the routing table. It assumes that the softswitch addresses of each domain are redistributed into the IGP. Because only a handful of addresses are expected per media gateway, this is a minor addition to the IGP. However, it does require the IGP to be clean from other nonessential addresses.

Unicast RPF Checking (uRPF)

The basis of this feature is to make sure the router does not forward traffic from unknown sources. The feature originates from IP Multicast, where it is used to verify that multicast packets are received from an interface that can be used to reach the source.

By installing PSTN-only routes in the telephony edge router's routing table, this feature prohibits forwarding traffic to any other destination. In essence, this creates a closed user group.

Packet Filtering

Routers can be configured to discard packets with the use of access lists. These are typically applied on peering links to disallow incoming packets to have internal (spoofed) destination addresses or to protect network operations center (NOC) and internal servers. Access lists can help prevent DoS attacks and are also an important tool when locking down attacks.

DoS Tracking

Any IP network should have procedures in place to track and lock down attacks. Further tools like the Cisco IOS DoS Tracker can be used to backtrack an attack to its source.

APPENDIX B. NETWORK SECURITY DISCUSSION

DoS Attacks

DoS attacks are reported increasingly in IP-based networks today. The anonymity and flexibility of IP and the Internet allow an attacker to remotely degrade router performance or even bring down end points including media gateways and IP infrastructure links or nodes that will impact a large number of users. Network-based DoS attacks flood the network with unwanted traffic, using up available bandwidth, CPU, memory, or some combination of these.

Some of these attacks may be targeted directly toward the router while others are targeted toward the transit of other hosts or servers (such as Domain Name System [DNS], Network Time Protocol [NTP], and Web servers) to a router. This unwanted traffic must be removed as close to the source as possible, because the available bandwidth typically decreases as the traffic moves closer to the destination. This means that tracking the offending source is important as well as applying access lists to protect the service or node under attack.

To avoid disruption of service or dropped calls, end systems as well as intermediate devices need to be strongly protected from unauthorized access and malicious attacks. If access lists are applied, policers can be based on source and destination addresses, TCP and UDP ports, protocol type, Internet Control Message Protocol (ICMP), etc.

The proposed solutions can be made more rigorous by deploying various security or protection mechanisms in media gateways as well as in telephony edge routers. For example, ingress filtering and rate limiters can be deployed at media gateways. Furthermore, BGP communities can be used to mark telephony addresses and then restrict access between end systems and force the core network and telephony edge routers to rate limit all traffic from nontelephony end systems destined to telephony end systems. This is absolutely necessary if the telephony edge routers are used for other purposes in parallel. The following security measures can help protect a network:

- The stub links should not be distributed into the service provider's backbone network
- No default routing is allowed toward telephony edge routers
- In some cases, the access control list has to be extended with rate limiters (maximum bandwidth and burst size)
- To protect against attacks toward the router, the route processor should be protected via receive ACLs.

Logging is important to track and detect intrusion. When ACLs are applied, all traffic that does not pass the match conditions can be logged. Scripts can parse the syslog files in real time and send an alarm to the NOC. There are other automated response systems that use known network behavior to trigger alarms.

IP Address Spoofing

Spoofing allows an intruder to pass IP packets to a destination as genuine traffic. To protect media gateways from spoofing attacks, uRPF is applied on all interfaces of telephony edge routers. The uRPF check performs a route-table lookup on an IP packet's source address, and checks the incoming interface. The router determines whether the packet is arriving from a path that the sender would use to reach the destination. If the packet is from a valid path, the router forwards the packet to the destination address. If it is not from a valid path, the router discards the packet. It is also possible to log the packet for follow up. This works very well for interfaces that connect to well-defined networks. It can also work for peers (over private connections or over an Internet exchange). The uRPF check function has to be applied on all border interfaces. To deal with dual-homed connections, it is necessary to apply uRPF check on active paths (paths that are in the forwarding table) as well as on the feasible paths (paths that are explicitly specified or resident in the routing table but not selected for forwarding).

Protection from Password and Intrusion Attacks

Password attacks usually refer to repeated attempts to identify a user account or password. Password attacks can be implemented using several different methods including brute force attacks. Brute force attacks are performed using a program that runs across the network and attempts to log into a shared resource, such as a server. An attacker who successfully gains access to a resource has the same rights as the user whose account has been compromised to gain access to that resource. If this account has sufficient privileges, the attacker can create a "back door" for future access, without concern for any status and password changes to the compromised account. To manage intrusion attacks, the host device needs to be proactive in addressing TCP/UDP port scans while maintaining detailed logs of login attempts and port scans. Security of the logs is essential so that a potential attacker cannot access them. The routers can also be configured to only allow access from certain NOC addresses.

Central Authentication

Administrative privileges can be granted based on user identity and password. User access control determines which administrators can access the device, regulates the level of access, and reduces administrative effort. Remote databases include systems such as RADIUS or TACACS+. In addition, one-time passwords are recommended for increased security.

Protecting Signaling, Charging, and Accounting

A malicious user could monitor and capture the packets that set up a call. By doing this, they could manipulate fields in the data stream and make VoIP calls without using VoIP, or they could initiate many expensive calls and make a local switch believe they are originating from a legitimate user. IP Security (IPSec) and Internet Key Exchange (IKE) can be used for signaling protection and even for payload protection if site-to-site protection is acceptable. It is, however, also necessary to protect operations and maintenance, charging, and billing information. Most of the existing signaling protocols (even SS7 for the circuit-switched network) lack authentication and contain vulnerabilities because they were not designed for connection to ubiquitous, global, untrusted networks, with the possibility of malicious attacks. To secure transport of SS7 signaling over an IP network, the IETF community has specified ways of transporting SS7 over IPSec and IKE. The IPSec standard addresses the authentication of peers, integrity of user data transport, and confidentiality of user data.

Securing Remote Node Access

Sometimes device management has to take place remotely and access to that device can only occur via a public network. In that case, encrypting protocols such as SSH Protocol may be required to secure management traffic. SSH Protocol is becoming a well-known standard for all remote command-line configurations and file transfers. If Web-based management is used, the Secure Sockets Layer (SSL) or Transport Layer Security (TLS) helps to secure HTTP traffic. SNMP is used to discover, monitor, and configure network devices. The secure implementation of SNMP Version 3 is essential because it helps ensure confidential and authenticated communications with reliable integrity.

Authenticating Routing

Securing your IGP and BGP will protect you from not only malicious attacks, but also accidental misconfigurations. The friendly nature of routing dictates that any router with coordinated configuration parameters (network mask, hello interval, dead interval, etc.) can participate in a given network. Because of this default behavior, any number of accidental factors (misconfigurations, lab machines, test setups, etc.) has the potential to adversely affect routing.

In OSPF, the security can be increased when authentication is enabled. OSPF authentication can be either none, simple, or MD5. With simple authentication, the password goes in plaintext over the network. Thus, anyone with a protocol analyzer on the OSPF network segment can pull the OSPF password, and the attacker would be one step closer to compromising your OSPF environment. With MD5 authentication, the key does not pass over the network. MD5 should be considered the most secure OSPF authentication mode.

In the IS-IS Protocol, protocol exchanges can be authenticated so that only trusted routers participate in the authentication server’s routing. This can be based on either simple authentication using a text password that is included in the transmitted packet, or Hash-Based Message Authentication Code (HMAC)-MD5 authentication that uses an iterated cryptographic hash function.

In BGP, all BGP exchanges can be authenticated to allow only trusted routers to participate in the authentication server’s routing. Authentication is done using an MD5 algorithm that creates an encoded checksum that is included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet’s MD5 checksum.

APPENDIX C. LATENCY DISCUSSION

Latency

Latency or delay is the time it takes a packet to make its way through a network end to end. A generally accepted end-to-end voice delay should be between 100–150 ms (from mouth to ear) for toll-quality phone calls. To keep delays below 150 ms, consider the following latency factors:

- **Packetization delay**—One source of latency is the packetization delay for voice services. This type of delay is caused by the amount of time it takes to fill a packet with data. Generally, the larger the packet size, the greater amount of time it takes to fill it. Packetization delay is governed by the codec standard being used. A codec takes an audio sample and digitizes it, often compressing it as part of the process. A codec may also take a compressed digital signal and return it to an analog audio sample. Normally, packetization delay does not exceed 30 ms.
- **Serialization delay**—Another source of delay is the time it takes to serialize the digital data onto the physical links. This delay is inversely proportional to the link speed.

  Serialization_delay = Frame_size / Link_speed

Serialization delay is considered negligible at links speeds above Synchronous Transport Module Level 1 (STM-1). A 1500-byte packet is clocked at STM-1 rate in 80 _s, at STM-16 rate in 5 _s, and STM-64 in 1.25 _s. Although this delay is unavoidable, keeping the number of intervening links small and using high-speed interfaces reduces the overall latency.

- **Propagation delay**—It takes time for an electrical signal to traverse the length of a conductor. Propagation delay is constrained by the speed of light in a medium and for optical fiber is around 5 ms per 1000 kilometers (km). This becomes an issue specifically when the signal travels a great distance. More precisely:

  *Propagation delay = Circuit km / 299,300 km * 0.6.*

  Propagation delay can vary as network topology changes when a link fails, or when an underlying network reroutes its circuit paths.

  Propagation delay also impacts echo. Echoes happen when callers hear their own audio delayed by as little as 25–30 ms. In this case, it is the round-trip delay that is important. In other words, the time it takes for a caller’s transmit (spoken) audio to travel to the receiving caller and return to the caller’s phone receiver. This means that as in circuit switching, an international call requires installment of echo cancellers.
- **Scheduling and queuing delay**—This is the amount of time a packet remains buffered in a network element while it awaits transmission. Network traffic loads result in variable queuing delays. The amount of buffer that a queue uses is usually a configurable parameter, and the smaller the better when it comes to latency. However, this delay is also based on the amount of traffic the element is trying to pass through a given link, and therefore it increases with the network load. Hence, if the queue used for engine telephony is not serviced quickly enough and the queue is allowed to grow too large, the result is greater latency.
- **Switching delay**—Switching or processing delay is the time difference between receiving a packet on an incoming router interface and the de-queuing delay of the packet in the scheduler of its outbound interface. Switching delays, however, are usually negligible in today’s high-performance routers—and are typically less than a tenth of a millisecond per packet.

Consequently, the total network delay in an IP network can be calculated as follows:

Total_network_delay = *Packetization_delay* + *Serialization_delay* + *Propagation_delay* + *Queuin_delay* + *Switching_delay*

In a well-designed network, queuing delay should be zero. That is apart from the transient congestion that may occur on any network. Among the delay components described previously, the propagation delay is the major source of delay by several orders of magnitude.

Jitter

Network-delay jitter characterizes the variation of network delay; it is generally computed as the variation of the delay for two consecutive packets. Jitter is caused by the variation in the delay components listed previously (such as propagation delay, switching delay, and scheduling delay). Jitter buffers are used to remove delay variation by turning variable network delays into constant delays at the destination end systems.

Jitter usually occur for the following reasons:

- Due to contention for resources between packets from multiple voice calls.
- Due to contention for resources between a real-time packet and a nonreal-time packet, which may cause variations in queuing delay.
- Due to equal-cost-balanced traffic taking physically different or longer paths.

This jitter must be bounded and minimized by the network to support real-time communication. Consequently, in networks that are engineered to support low-delay services such as telephony, it is important that they are also engineered for low jitter. Large buffer size will increase the overall delay. On the other hand, small buffer size will minimize the overall delay but will limit the jitter that can be handled. IP backbones that are engineered to support high-quality services typically budget for 5–10 ms of jitter in the backbone network; assuming 10 backbone hops, this gives a jitter budget per hop of 500–1000 ms. If the output queues used for the voice traffic are kept relatively small, statistically, jitter becomes a less of an issue. This is of course given that the voice queue is serviced in proportion to the amount of traffic it is expected to schedule for transmission.

REFERENCES

- Cisco 12000 Series Router architecture:

  http://www.cisco.com/go/12000
- Ericsson ENGINE Integral architecture:

  http://www.ericsson.com/products/main/ENGINE_Integral_hpsol.shtml
- ISP Essentials, Green & Smith, ISBN 1587050412:

  http://www.ciscopress.com
- Routing design concepts—please refer to the recommended reading list
- [ITU-T Y.1541] Network performance objectives for IP-based services
- [ITU-T G.114] One-way transmission time

RECOMMENDED READING—BOOKS

Title	Author	ISBN
BGP Design and Implementation (Cisco Press®)	Zhang, Bartell	1587051095
ISIS Network Design Solutions (Cisco Press)	Abe Martey	1578702208
OSPF Network Design Solutions (Cisco Press)	Tom Thomas	1587050323
Traffic Engineering with MPLS (Cisco Press)	Osborne, Simha	1587050315
Troubleshooting IP Routing Protocols (CCIE® Professional Development Series, Cisco Press)	Aziz, Liu, Martey, Shamim	1587050196
Voice-Enabling the Data Network: H.323, MGCP, SIP, QoS, SLAs, and Security (Cisco Press)	James Durkin	1587050145
Engineering a multiservice IP backbone to support tight SLAs (Computer Networks, Volume 40, Issue 1, September 2002, pp. 131_148)	Clarence Filsfils, John Evans	

RECOMMENDED READING—INTERNET RESOURCES

Name	Description	Link
Tutorial: Deploying Tight-SLA Services on an Internet Backbone - ISIS Fast Convergence and Differentiated Services Design	Verified QoS Design Considerations	http://www.nanog.org/mtg-0206/filsfils.html
Designing Large-Scale IP Internetworks	Routing Protocol Design Concepts	http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd2003.htm
Cisco OSPF Design Guide	Design Concepts	http://www.cisco.com/warp/public/104/1.pdf

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

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the
Cisco Website at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe