



Cisco Cable IP Solutions for High-Availability Networks



TABLE OF CONTENTS


| | |
|----------------------------------------------------------------------------------------------|----|
| EXECUTIVE SUMMARY | 3 |
| 1 DEFINING AND CALCULATING AVAILABILITY | 5 |
| Calculating Availability | 7 |
| 2 CABLE VOICE OVER IP: CABLELABS PACKETCABLE AVAILABILITY REQUIREMENTS AND ANALYSIS | 8 |
| 3 CISCO® CABLE SOLUTIONS: PACKETCABLE VOIP | 11 |
| 4 CISCO PACKETCABLE VOIP SOLUTION AVAILABILITY ANALYSIS | 13 |
| Reference Architecture | 14 |
| High-Availability Features of Solution Components and Assumptions | 15 |
| Cisco uBR10012 Series Universal Broadband Router | 15 |
| Cisco 12000 Series Router | 15 |
| Cisco Catalyst® 6500 Series Switch | 15 |
| Cisco BTS 10200 Softswitch | 15 |
| Cisco MGX® 8800 Series Gateway | 15 |
| Impact of Unplanned Failure Scenarios and Planned Downtime | 16 |
| Cisco uBR10012 Series Universal Broadband Router | 16 |
| Cisco 1200 Series Router | 16 |
| Cisco Catalyst 6500 Series Switch | 17 |
| Cisco BTS 10200 Softswitch and Provisioning Servers | 17 |
| Cisco MGX 8800 Series Gateway | 17 |
| Planned Downtime | 18 |
| Methodology and Assumptions | 19 |
| Summary of Results | 21 |
| 5 CONCLUSION | 23 |



EXECUTIVE SUMMARY

Cable operators are migrating toward converged network architectures to achieve minimized capital and operational expenditures and realize higher returns on invested capital. As operators converge critical services such as voice over IP (VoIP) to common network architectures, the requirements on the IP network (Figure 1) become much more stringent (compared to requirements on best-effort data).

Figure 1
Network Requirements on PSTN-Grade Voice




As shown in Figure 1, networkwide availability becomes a critical aspect of offering advanced services such as VoIP. A highly available network is required to meet service-quality goals, provide a competitive service offering and a rich user experience, and minimize ongoing operating expenses.

To establish an industry benchmark, CableLabs has specified a set of availability metrics for cable VoIP services based on the PacketCable model. These metrics are derived from analysis of existing circuit-switched telephony models, with the goal of meeting or exceeding traditional voice-service availability using VoIP technology.

Cisco Systems[®] is the market leader in highly available IP networking solutions. In the cable multiple system operator (MSO) space, Cisco[®] provides a highly integrated, PacketCable-compliant VoIP solution. The market leadership of Cisco in IP high availability is evident in software and hardware redundancy and resiliency features such as Non Stop Forwarding (NSF); Stateful Switchover (SSO), routing optimization technologies for fast convergence; fiber-transport technologies with rapid failure recovery (such as Dynamic Packet Transport (DPT), Resilient Packet Ring (RPR), Synchronous Optical Network (SONET), and wavelength protection switching); and other innovative availability solutions.

Additionally, Cisco has extensive experience in partnering and working with the world's largest service providers to develop operational best practices to further increase service availability.

Detailed analysis of Cisco's PacketCable VoIP solution indicates that Cisco can meet and exceed PacketCable VoIP availability requirements today. The results of this analysis show that the commitment by Cisco to cable high availability has resulted in cable VoIP solutions capable of exceeding traditional circuit-switched voice-service availability metrics. Cisco continues to add high-availability features in its network elements for the PacketCable VoIP solution to increase VoIP availability in cable networks.



1. Defining and Calculating Availability



Reliability, resiliency, and availability are sometimes used interchangeably. However, although all three terms are related to the concept of high availability, it is important to note the differences in terminology. Reliability is the probability that a system will not fail during a specified period of time. Resiliency is the ability of a system to recover to its normal operating form after a failure or an outage. Availability is the ratio of time that a service is available to total time.

Availability can be expressed as mean time between failure (MTBF) and mean time to repair (MTTR), and expressed in mathematical terms as:

$$\text{Availability} = \text{MTBF}/(\text{MTBF}+\text{MTTR})$$

MTBF is tied to the reliability of the system, while MTTR and resiliency are closely related. Thus system availability increases as the reliability and/or resiliency of the system is increased. Availability is typically expressed in percentage of time the system is available or in downtime per year. The two methods of expressing availability are equivalent and related as shown in Table 1.

Table 1 Methods of Expressing Availability

| Number of 9s | Availability | Downtime per Year | Types of System |
|--------------|--------------|------------------------|------------------------------------------|
| 1 | 90.0000% | 36 days, 12 hours | |
| 2 | 99.0000% | 87 hours, 36 minutes | Fax machines, printers |
| 3 | 99.9000% | 8 hours, 46 minutes | Dial ISPs, non-critical business systems |
| 4 | 99.9900% | 52 minutes, 33 seconds | Data centers |
| 5 | 99.9990% | 5 minutes, 15 seconds | Redundant storage arrays |
| 6 | 99.9999% | 31.5 seconds | Aviation and military defense |

This definition for availability is good for a simplex system (a system comprising of one element). However, in a network that consists of multiple trunks and routers, most failures are partial failures. Service disruptions resulting from such failures typically affect a subset of customers, while service for others is uninterrupted. Further, even within a network element only one component may fail, thus only disrupting service for a subset of users serviced by the network element. Hence availability is defined with respect to a single customer of the network. Therefore, to compute availability, it is necessary only to consider the components along the path needed to provide service to a single customer.

Another important metric is defects per million (DPM). DPM is defined as the number of defects that an end customer would experience per million events. This metric is adaptable to the application being run on the network and relates directly to the service availability as seen by the end user. In the telephony space there are two critical metrics for service availability: DPM for calls dropped and DPM for ineffective attempts.

DPM for calls dropped refers to the number of calls per million that are dropped while in progress.

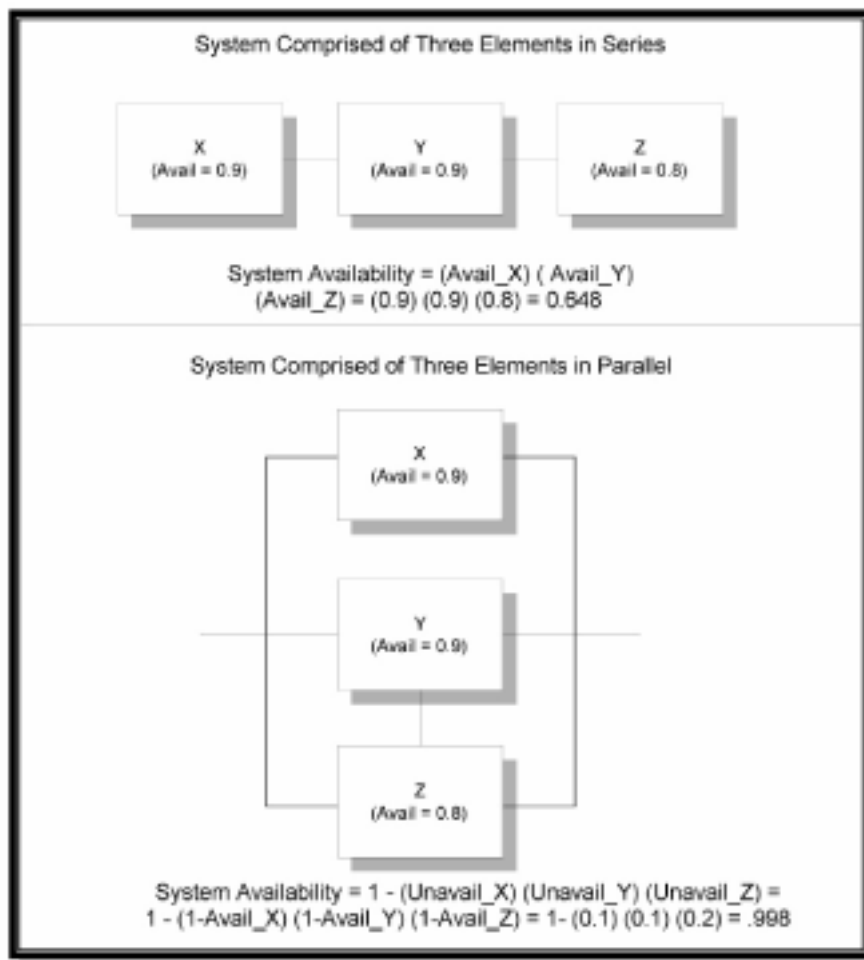
DPM for ineffective attempts refers to the number of calls per million that cannot be set up due to a failure in the signaling path.

Calculating Availability


The common availability definition of $MTBF/(MTBF+MTTR)$ does not fully address complex systems. Real-world networks are a complex mixture of serial and parallel network elements. If components are combined in series, the overall network relies on the availability of all components and the total system availability can be much lower than the availability of the weakest component. However, if components are combined in parallel, with redundancy provided by the parallel components, the system availability can be higher than that of the most available component.

For systems comprised of elements in series, the overall system availability is the product of the individual element availability values. For systems comprised of elements in parallel, the availability of the system is one minus the product of the individual component unavailability values (where unavailability = 1 – availability). Figure 2 illustrates this concept.

Figure 2
System Availability Elements



Using these basic principles of availability modeling, it is possible to analyze an end-to-end cable network and determine whether or not it meets the PacketCable specifications. Each network element can be modeled as a serial or parallel component. In addition, parallel elements can be used to model redundancy in conjunction with other parameters such as switchover time, and active and standby coverage. Details behind these parameters are beyond the scope of this document and are described in a separate paper published by Cisco. A copy of this paper can be made available under Non-Disclosure Agreement (NDA). Please contact your Cisco sales representative for further information.



2. Cable Voice over IP: CableLabs PacketCable Availability Requirements and Analysis




Founded in 1988, CableLabs is focused on developing specifications and industry standards for various cable telecommunications services and ensuring that these standards meet cable operator technical and business objectives. An initiative within CableLabs known as PacketCable is responsible for developing specifications for various real-time services over two-way cable plant. Cisco has been an important contributor to CableLab's PacketCable initiative since the initiative's inception through significant contributions to PacketCable architectures and requirements.

The PacketCable 1.x series specifications define an end-to-end architecture for cable VoIP services. In the area of availability, a PacketCable 1.1 Technical Report addresses VoIP availability and reliability models (VoIP Availability and Reliability Model for the PacketCable Architecture Technical Report, Technical Report Number pkt-tr-voipar- v01-001128). These specifications define the end-to-end architecture for PacketCable VoIP and the associated availability metrics derived through analysis of existing circuit-switched voice architectures.

The PacketCable VoIP availability model (Figure 3) is derived from analyzing existing Telcordia (formerly Bellcore) GR series specifications for traditional circuit-switched voice telecommunications systems and deriving equivalent availability requirements in the PacketCable VoIP domain.

Figure 3
PacketCable Availability Model



This model specifies several availability metrics for cable VoIP:

- End-to-end availability (99.94 percent)—This availability metric was derived from Telcordia availability specifications for end-to-end traditional circuit-switched telecommunications networks. This metric also implies total downtime for the end-to-end cable VoIP network.

- Subsystem availability (for instance, 99.975 percent for managed IP network)—These availability metrics were derived through analysis of Telcordia subsystem availability metrics in the circuit-switched reference network model. This metric also implies per-subsystem downtime budgets.
- End-to-end service availability metrics of less than 125 DPM for calls dropped and less than 500 DPM for ineffective attempts—These metrics are extremely important because they map directly to end-user service experience.

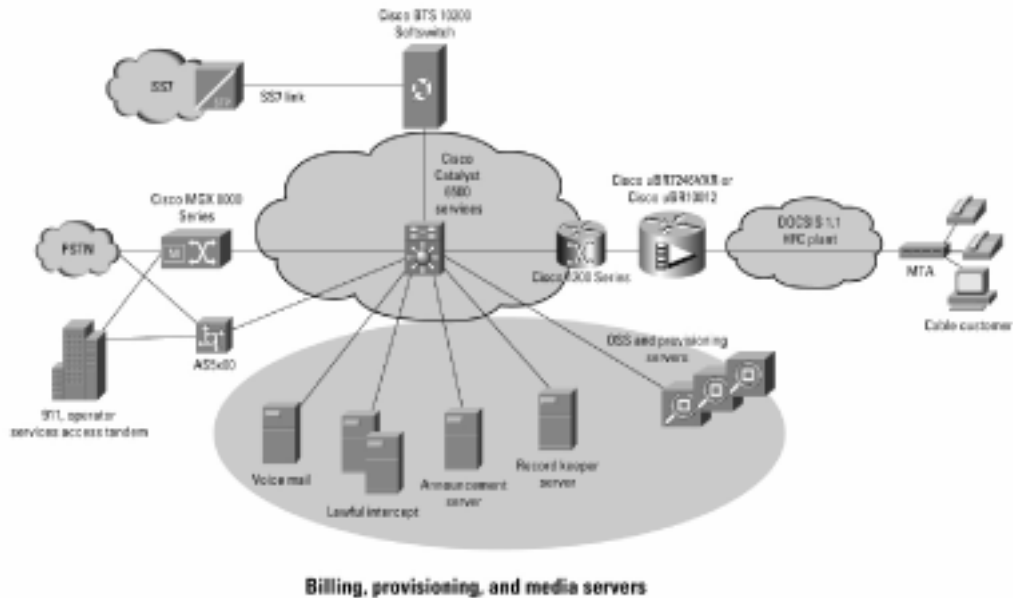
Further analysis of availability requirements in this PacketCable Technical Report is outside the scope of this paper. Additional details are contained within the technical report (VoIP Availability and Reliability Model for the PacketCable Architecture Technical Report, Technical Report Number pkt-tr-voipar- v01-001128) at <http://www.packetcable.com/specifications/>. For a more detailed analysis of PacketCable requirements and implications for an end-to-end environment, refer to a separate paper on this subject published by Cisco. A copy of this paper can be made available under Non-Disclosure Agreement (NDA). Please contact your Cisco sales representative for further information.



3. Cisco Cable Solutions: PacketCable VoIP

Cisco PacketCable VoIP solutions allow cable operators to deploy VoIP services based on CableLab's PacketCable specifications, helping ensure compliance to industry standards, interoperability, and investment protection. The Cisco solution combines world-class Cisco products with Cisco partner products to provide a complete carrier-class cable VoIP service. Figure 4 illustrates this solution.

Figure 4
Cisco PacketCable Solution



The Cisco PacketCable solution is comprised of the following components:

- The Cisco uBR7246VXR and uBR10012 universal broadband routers—These routers combine the functions of a router and a cable modem termination system (CMTS) in an integrated platform supporting Data over Cable Service Interface Specification (DOCSIS and EuroDOCSIS configurations, and are qualified for DOCSIS1.0, 1.0+, and 1.1). The CMTS is actively involved in such PacketCable functions as network security, provisioning, dynamic quality of service (DQoS), and event messaging necessary for rated billing.
- Cisco 12000 Series routers and Cisco Catalyst 6500 Series and Cisco 7600 Series—These carrier-class core and edge routing and switching platforms provide a highly scalable and flexible solution for building IP backbone networks.
- The Cisco BTS 10200 Softswitch—This component provides the call-management-server function of the telephony solution. It provides call-control intelligence for local, long-distance, toll-free, 900 and 976 voice services, and the subscriber features that previously required implementation of a Class 5 switch-based architecture. Additionally, the Cisco BTS 10200 serves as an interface to enhanced service and application platforms such as unified messaging, conferencing, and many other revenue-generating services.
- The Cisco MGX 8000 Series and the Cisco AS5000 Series media and trunking gateways—These high-capacity carrier-class voice gateways offer standards-based support for VoIP.
- Operations support systems (OSSs) and provisioning systems—These systems provide provisioning, monitoring, and management of PacketCable functions within the network.
- Third-party partner solutions—These include media terminal adapters (MTAs) and voice-mail, announcement, record-keeper, and lawful-intercept servers.



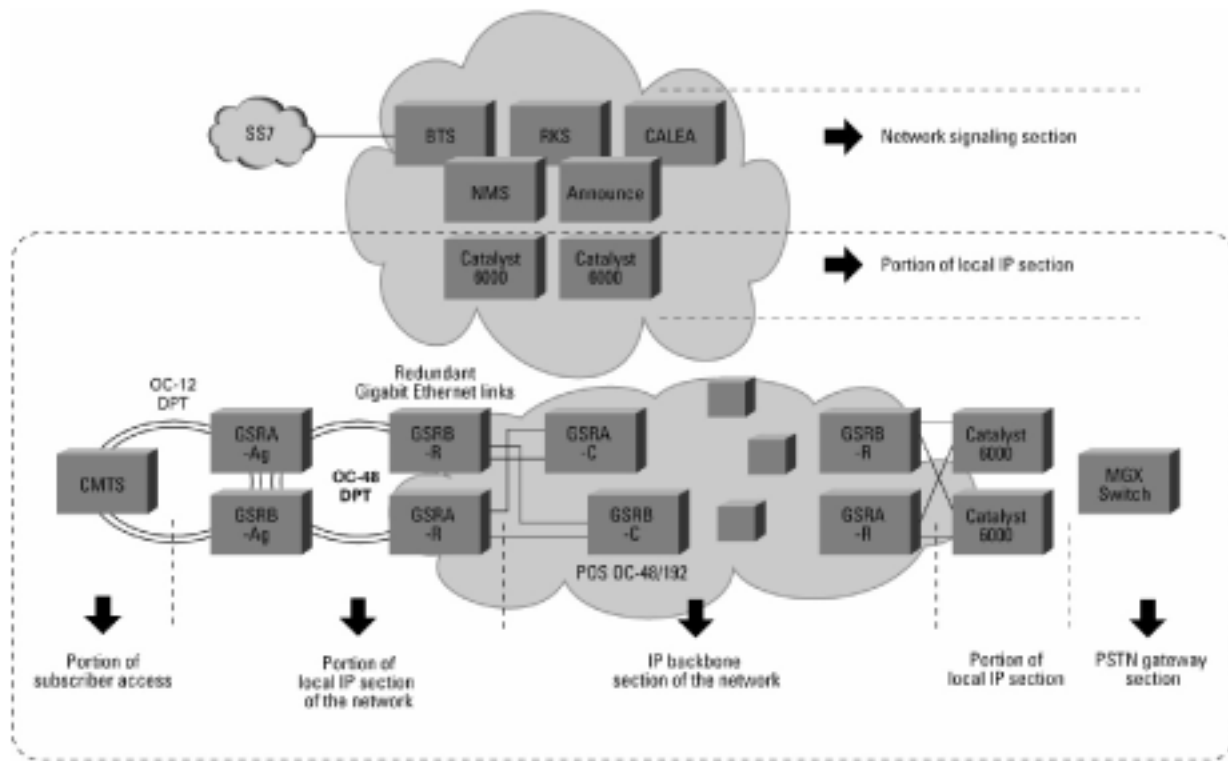
4. Cisco PacketCable VoIP Solution Availability Analysis



Reference Architecture

A reference architecture is used in this paper to analyze service availability of the Cisco PacketCable VoIP solution. This architecture is based on the existing Cisco PacketCable VoIP solution as discussed in the previous section. More specifically, the reference architecture specifies the types of interfaces between network devices, hardware and software redundancy features built into network devices, and other parameters such as routing protocol design. The goal of the analysis is to calculate availability for the VoIP reference architecture and compare it to PacketCable VoIP reference architecture (Figure 5).

Figure 5
Cisco VoIP Reference Architecture



This architecture consists of a Cisco uBR10012 CMTS at the edge of the cable operator network connected through an OC-12 Dynamic Packet Transport (DPT) to redundant Cisco 12000 Series aggregation routers (portrayed as GSRs in Figure 5). The aggregation routers are connected through an OC-48 DPT ring to redundant regional Cisco 12000 Series routers. These regional routers connect to redundant Cisco Catalyst® 6500 Series switches via Gigabit Ethernet for connectivity to the Cisco BTS 10200 Softswitch and PacketCable billing, media, and provisioning servers. The regional routers connect to a core IP-transport network consisting of additional Cisco 12000 Series routers via redundant packet-over-SONET (POS) links. At the other edge of the core IP-transport network, a pair of redundant Cisco 12000 Series routers is connected to Cisco Catalyst 6500 Series switches that are connected to a Cisco MGX 8800 media gateway. The MTA and the hybrid fiber-coaxial (HFC) plant are not included in this analysis. These components are typically nonredundant and calculation of their availability is outside the scope of this paper. However, reference can be made to PacketCable availability budgets allocated to these components to ensure that complete end-to-end availability requirements are satisfied.

High-Availability Features of Solution Components and Assumptions

This paper now examines the redundancy and high-availability features of this reference Cisco PacketCable VoIP architecture. This section will describe assumptions made for solution components with respect to hardware and software redundancy, high-availability features, and routing-convergence optimization.

The Cisco uBR10012 CMTS, Cisco 12000 Series, Cisco Catalyst 6500 Series switches, Cisco BTS 10200 Softswitch, and Cisco MGX 8800 Series trunking gateway all provide numerous high-availability features, including hardware redundancy and software features for fast switchover and recovery in case of failure.

In the reference architecture the following assumptions are made for each of these components with respect to high-availability features and design:

Cisco uBR10012 Universal Broadband Router

- Redundant route processor modules, power supplies, control cards, and cooling.
- Redundant cable DOCSIS line card in N+1 mode.
- Route Processor Redundancy Plus (RPR+) and DOCSIS SSO support in Cisco IOS[®] Software for fast route processor failover without DOCSIS line card reboot.
- Routing optimization for fast routing convergence after route processor failover.
- OC-12 and OC-48 DPT uplink providing sub-50 milliseconds (ms) recovery for fiber cut or ring failure.

Cisco 12000 Series Router

- Redundant route-processor modules, power supplies, control cards, and cooling.
- Cisco IOS NSF and SSO enabled for POS and DPT line cards (immediate forwarding recovery from route-processor failure without routing-adjacency loss).
- Routing optimization for fast routing convergence due to link failures, topology changes, or route-processor failover (routing convergence within 2-3 seconds).
- DPT provides sub-50 ms recovery for fiber cut or ring failure.

Cisco Catalyst 6500 Series Switch

- Redundant supervisor-engine modules, power supplies, switch-fabric cards, and cooling.
- Cisco IOS Software for RPR+ for Ethernet line cards.
- Dual-homing of PacketCable control servers to dual Cisco Catalyst 6500 Series switches.
- Dual-homing of Cisco MGX 8800 Series trunking gateway to dual Cisco Catalyst 6500 Series switches.

Cisco BTS 10200 Softswitch

- Any failure in the Cisco BTS 10200 Softswitch will not affect calls in progress. Failure in the Cisco BTS 10200 (or any other component in the call-management path for that matter) will only influence the ineffective-attempts metric.
- The Cisco BTS 10200 is deployed in an active standby mode and consists of a redundant pair of Sun servers connected to a pair of redundant Cisco Catalyst 2924 switches. If there is a failure on either of the servers or the switches, or the provisioning server, the outage duration is approximately seven seconds.
- ICMP Router Discovery Protocol (IRDP) is configured between the Sun servers and the Cisco Catalyst 6500 Series switches (the Cisco Catalyst 2924 switches are in simple Layer 2 switching mode and do not perform IP routing). The recovery time for IRDP is approximately seven seconds.

Cisco MGX 8800 Series Gateway

- Cisco Voice Internetworking Service Module (VISM) failures are recovered through Cisco BTS 10200 Softswitch rerouting of calls to an alternative gateway in approximately four seconds.
- OC-3 link to the Cisco MGX 8800 Service Resource Module Enhanced (SRM-E) for high availability.
- Permanent virtual circuit (PVC) failover mechanism between RPMs for faster RPM recovery.

Impact of Unplanned Failure Scenarios and Planned Downtime

This section examines unplanned failure scenarios and planned downtime on different products and their impact on availability. Each product has a different set of high-availability features that will be released over time, and these features will help reduce the outage time in the event of a component failure.

Cisco uBR10012 Universal Broadband Router

- Route processor (performance routing engine (PRE) on Cisco uBR10012) failover—With RPR+ and DOCSIS SSO, the Cisco uBR10012 can rapidly fail over from the active route processor to the standby processor without the reloading of the cable line cards. However, even though the cable line cards are not reset, the new active route processor needs to perform certain recovery procedures in order for cable line card traffic-flow to resume. A Cisco implementation provides priority-recovery procedures for those modems carrying voice, providing more rapid recovery of voice services.
- Cable line card failover—With N+1 cable line card redundancy, a standby cable line card can rapidly recover a failed line card. Moreover, once the original active line card is restored, mechanisms exist for the switchback to occur within 2-3 seconds. Switchover coverage (percentage of events where failure detection succeeds and thus system is recovered through automatic failover) is conservatively assumed at 98 percent. Therefore, two percent of RF failures will require manual recovery averaging 15 minutes in duration.
- DPT uplink module failure—With routing optimization, routing protocols can converge rapidly (2-3 seconds) to restore packet flow and voice service.
- Fiber cut or other ring failure— DPT's inherent protection mechanisms provide sub-50 ms protection in these failure scenarios.

Cisco 12000 Series Router

- Route-processor failover—For high availability, each Cisco 12000 Series is deployed with dual route processors and Cisco IOS Software NSF and SSO capabilities. With NSF/SSO, the Cisco 12000 Series can fail over from the active to standby route processor almost immediately while continuing to forward packets through the line cards. This feature does require that adjacent routers be NSF-aware (at a minimum) or NSF-capable and thus be able to continue forwarding packets to the router undergoing failover. The Cisco uBR10012 and the Cisco Catalyst 6500 Series will support this feature set in subsequent software releases, resulting in zero packet loss during a Cisco 12000 Series route processor failover event. Currently, on the router connected to a Cisco uBR10012 or Cisco Catalyst 6500 Series routing optimization can be employed to target reducing packet-forwarding recovery time to less than three seconds. On router-to-router links, packet-forwarding recovery time is essentially zero because full NSF/SSO capabilities are available.
- Line-card failover: OC-12 and OC-48 DPT line card failures will be recovered through Layer 3 routing-protocol mechanisms. With routing optimization, the recovery time can be reduced to approximately 2.5 seconds.
- Fiber cut or other ring failure: DPT's inherent protection mechanisms provide sub-50 ms protection in these failure scenarios.

Cisco Catalyst 6500 Series Switch

- Supervisor-engine failover—With support for RPR+ on the Cisco Catalyst 6500 Series switch, a supervisor-engine failover can occur without the need to reset the Gigabit Ethernet or Fast Ethernet line cards. Through route optimization, packet forwarding can recover from a supervisor-engine failover within approximately 2.5 seconds. Future Cisco IOS Software NSF/SSO support on this platform will enable immediate failover.
- Gigabit Ethernet failure—These links are used to connect to core Cisco 12000 Series routers and Cisco MGX 8800 trunking gateways. With route optimization, packet forwarding can recover from this failure within approximately 2.5 seconds.
- Failure of Gigabit Ethernet or Fast Ethernet uplink to control servers—The control servers are dual-homed to the Cisco Catalyst 6500 Series switches. The control servers will employ a mechanism through which they will detect a lost link and allow them to fail over to the redundant Cisco Catalyst 6500 Series within five seconds.

Cisco BTS 10200 Softswitch and Provisioning Servers

- Cisco BTS 10200 Softswitch outage—Any outage in the Cisco BTS 10200 or the provisioning servers should not cause existing calls to be dropped. Only new call attempts will be affected during the outage.
- Cisco BTS 10200 servers—The Cisco BTS 10200 redundant servers and the provisioning servers are dual-homed to the Cisco Catalyst 6500 Series switches for high availability. IRDP is used to recover from failures. Recovery time is about seven seconds.

Cisco MGX 8800 Series Gateway

- RPM failover—A PVC mechanism can be implemented between the redundant RPM cards, reducing the system downtime to less than two seconds.
- Cisco VISM failover—The Cisco MGX 8800 Series supports Cisco VISM failover. However, existing calls are dropped during the failover process. Cisco will deliver a future voice module for the Cisco MGX 8800 Series that will support stateful hot-standby failover. Meanwhile, the VoIP network can be designed to recover from trunking gateway failover in a more efficient fashion using the Cisco BTS 10200 Softswitch call-agent capabilities. Upon a Cisco VISM failure, the softswitch diverts all new calls to other Cisco VISMs, hence the user sees an outage equivalent of approximately four seconds.
- Cisco VISM back card—Using the Cisco VISM back card for T1 input can create a single point of failure. Higher redundancy and availability is achieved by using an OC-3 circuit into the SRME. That way, if there is a failure, network recovery is achieved in less than 250 ms and there is virtually no impact on voice calls.
- Processor Switch Module (PXM)—The Cisco MGX 8800 supports PXM redundancy and failover operation.

The Cisco MGX 8800 Series failure scenarios and their characterized impact are illustrated in Figure 6.

Figure 6
Effects of Unplanned Failure Scenarios and Planned Downtime for Cisco MGX

| | PXM (F&B) | RPM (F) | RPM (B) | Cisco VISM (F) | Cisco VISM (B) | SRM (F) | TOTAL |
|-------------------------------|-----------|-----------|-----------|----------------|----------------|---------|---------|
| Processor Switch Module (F&B) | <250 ms | NO EFFECT | | | | | <250 ms |
| RPM (F) | | <2 secs | NO EFFECT | | | | <2 secs |
| RPM (B) | | | <2 secs | NO EFFECT | | | <2 secs |
| Cisco VISM (F) | | | | 4 secs | NO EFFECT | | 4 secs |
| Cisco VISM (B) | | NO EFFECT | | | 4 secs | | 4 secs |
| SRM (F) | | NO EFFECT | | | | <250 ms | <250 ms |

Note: F = Front, B = Back

Planned Downtime

The failure scenarios examined previously were attributed to unplanned failures or unplanned downtime. In addition to unplanned downtime, planned downtime is extremely important when determining overall system availability. Planned downtime could result from a number of normal network activities, including software or hardware upgrades and capacity augmentation. To maximize service stability and system availability, these types of downtime need to be carefully planned and executed.

The availability analysis makes the following assumptions for downtime for software upgrades:


- Cisco uBR10012 software upgrades—Between 5 and 15 minutes of downtime. Fast Software Upgrade (FSU) will be delivered on this platform, bringing this downtime to less than 30 seconds in the future.
- Cisco 12000 Series Router software upgrades—Less than six seconds of downtime per router pair. This downtime is derived from the assumption that redundant routers exist at every point in the network and the upgrades are conducted one router at a time. Thus, through network redundancy and routing optimization, the impact from each software upgrade is less than three seconds.
- Cisco Catalyst 6500 Series Switch software upgrades—Approximately 30 seconds of downtime. Similar to the Cisco 12000 Series routers, the Cisco Catalyst 6500 Series switches are deployed in a redundant manner and thus downtime from software upgrades (performed one at a time) can be significantly reduced.
- Cisco MGX software upgrades—No effect on voice calls. If properly planned, a Cisco VISM can be upgraded without affecting existing calls. This is achieved through redistributing calls to one Cisco VISM, upgrading software on the other Cisco VISM, and restoring calls to the upgraded Cisco VISM. The same process described for the Cisco 12000 Series routers and Cisco Catalyst 6500 Series switches can be used to upgrade RPM software as well.

Methodology and Assumptions

The previous section examined each component of the Cisco PacketCable VoIP solution and characterized various failure scenarios and their impact on voice calls. This section will take a holistic approach to calculating the solution availability and discuss definitions, methodologies, and assumptions used to derive system availability.

Referring once again to the PacketCable availability model, the availability metrics for the IP portion of the network (that is, without including the HFC plant, MTA, and off-net circuit-switched voice network) are illustrated in Figure 7. CD represents calls dropped and IA represents ineffective attempts.

Figure 7
Availability Metrics for the IP Portion of the Network




For the IP portion of the network, the DPM metric for calls dropped is 65 and the metric for ineffective attempts is 275 DPM. These numbers are derived through analysis of non-IP portions of the network (cable modem or MTA, HFC plant, local exchange, and access network) and approximating their individual contributions to DPM calls dropped and DPM ineffective attempts and subtracting those from the overall budgets.

Calls dropped and ineffective attempts will be the metrics by which the Cisco PacketCable VoIP solution will be compared because these metrics directly correlate to the user experience.

A calls-dropped defect event will occur when an active voice call experiences an outage (silence) for more than three seconds. This assumption is based on the fact that in a VoIP environment, a network failure can cause active voice calls to go silent for a short period of time even though the call itself has not been torn down. Three seconds is chosen as a conservative threshold for how long a subscriber will wait for restored service before determining that the call has been abnormally “terminated” and terminating the call on purpose. Similarly, an ineffective-attempt defect event will occur when a subscriber wishes to make a valid call and cannot do so for a period of more than 30 seconds. This is the assumed threshold for the amount of time that the subscriber is willing to retry the call attempt. With these definitions now in place, the failure scenarios described in previous sections can be mapped to their contribution to calls dropped and ineffective attempts.

As noted previously, planned downtime can significantly contribute to DPM metrics. Because this type of activity is planned, it can and should be implemented in such manner as to minimize any service outage. Hence, network engineers and operations personnel should plan these types of downtime for periods of off-peak usage (for instance, 3 a.m.) when traffic load is significantly lower than average usage periods. Limiting planned downtime to these off-peak hours results in less DPM. The relative contribution of the planned and unplanned DPM (ineffective attempts) to the total DPM (ineffective attempts) will be the ratio of the mean call rate to the call rate during the planned downtime. For the purpose of this analysis, a conservative assumption of 50 percent is used. Thus the network will handle half as many calls during the planned downtime period as compared to the average call load on the network. Using this assumption, the availability calculations will factor down the DPM contributions of planned downtime by one-half as compared to unplanned downtime (Figure 8).

Figure 8
The Night Factor in Planned Downtime



One final assumption is made regarding availability and DPM calculations. Each component in the network is assumed to have two planned downtime events per year.

Summary of Results

Based on the reference Cisco PacketCable VoIP architecture and the assumptions discussed in this paper, a series of system-availability calculations are used to derive the resulting DPM metrics. The calculations are performed for each phase of high-availability feature enhancements in the Cisco PacketCable solution (Figure 9). The resulting DPM metrics are shown in Figure 10.

Figure 9
Phases of High-Availability Feature Enhancements

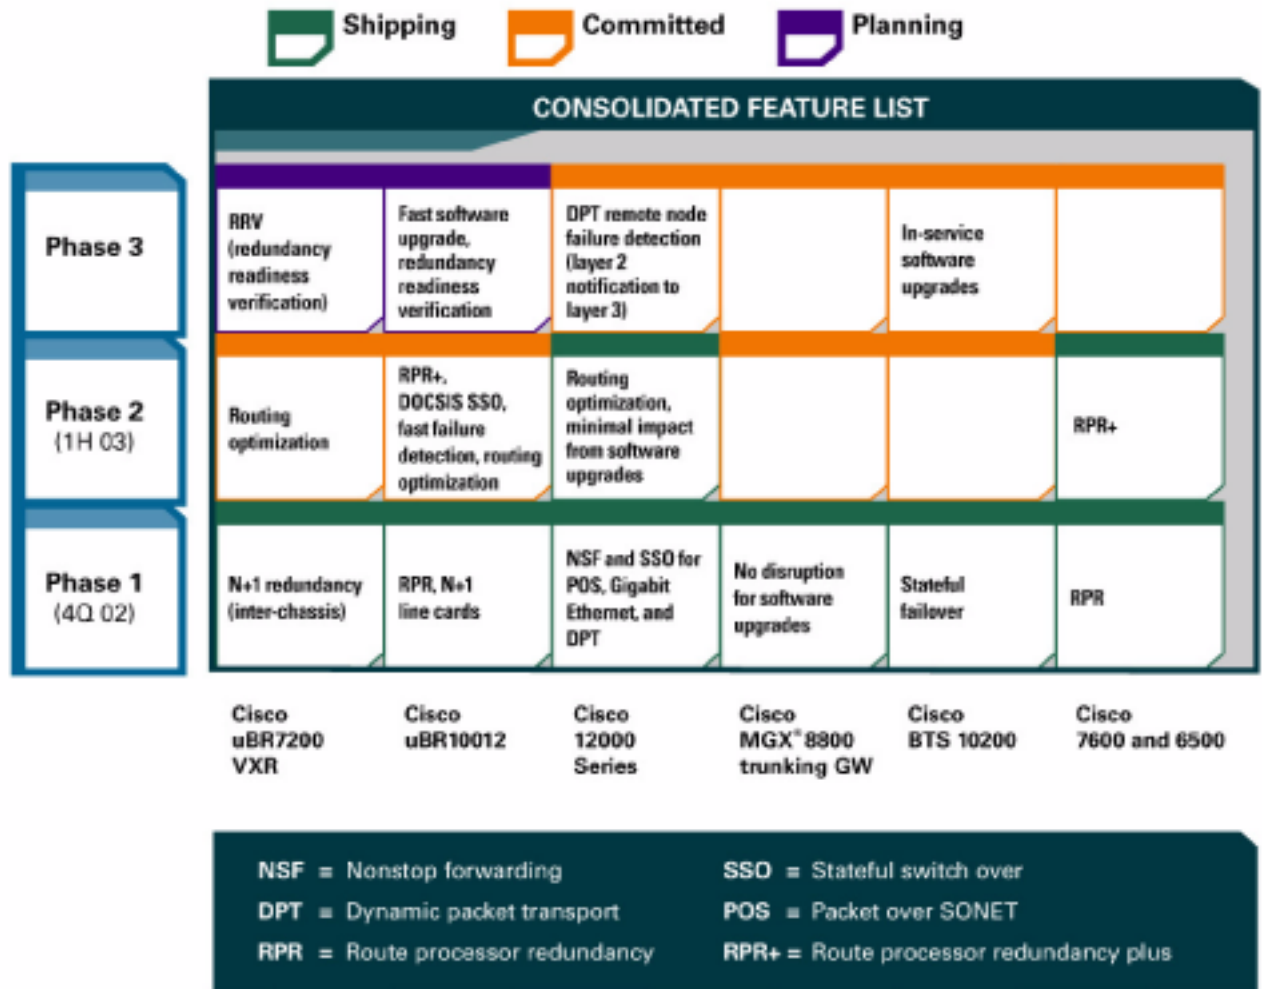



Figure 10
Summary of DPM Metrics



Note: Note: yellow bar denotes contribution from unplanned downtime, while gray is planned downtime contribution.

The results indicate that the Cisco reference solution does exceed PacketCable availability requirements. More specifically, PacketCable specifies 65 DPM for calls dropped and 275 DPM for ineffective attempts. The Cisco reference architecture clearly meets these DPM requirements (calls dropped in Phase 2, ineffective attempts in Phase 1).

The results also indicate that planned downtime can significantly contribute to service DPM. Thus, establishment of operational best practices and continual education of operations personnel are critical ingredients to increasing service availability.



5. Conclusion



Technological improvements at the component, system, transport, network, and routing level are enabling IP networks to exceed traditional circuit-switched networks in terms of resiliency and availability. Cisco is at the forefront of these enhancements and continues to lead the industry in extending the availability of IP networks.

PacketCable-based VoIP services provide cable operators with enticing service-bundling packages and incremental revenue streams from existing and future broadband subscribers. Cable operators are looking to use existing broadband IP assets in delivering a broad range of VoIP services. Consumer expectations of voice quality and availability will compel cable operators to implement highly available broadband IP networks in support of cable VoIP. Not surprisingly, PacketCable's availability requirements for VoIP services were derived from existing circuit-switched telephony availability requirements. Thus, a VoIP solution meeting or exceeding PacketCable requirements provides equivalent or better availability than existing circuit-switched telephony networks.

The reference architecture studied in this paper can exceed PacketCable availability requirements and provide extremely reliable voice services. Industry-leading Cisco technologies, coupled with broadband IP expertise and operational experience, enable cable operators to design, implement, and operate PacketCable compliant VoIP networks that meet or exceed the availability of traditional circuit-switched voice. Cisco VoIP solutions continue to evolve to address cable operator requirements for highly available networks today and into the future.



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.
Capital Tower
168 Robinson Road
#22-01 to #29-01
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 Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
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

All contents are Copyright © 1992–2003 Cisco Systems, Inc. All rights reserved. CCIP, CCSP, the Cisco Arrow logo, the Cisco *Powered* Network mark, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, MGX, MICA, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.
(0304R) RDA/JB 4711-06/03