

# VoIP Reliability: A Service Provider's Perspective

Carolyn R. Johnson, Yakov Kogan, Yonatan Levy, Farhad Saheban, and Percy Tarapore, AT&T Labs

## ABSTRACT

Voice over IP services offer important revenue-generating opportunities, as well as many technical challenges in providing high-quality services. Users have come to expect highly available telecommunications services with high-quality voice. Service providers need reliable high-performance networks to meet user expectations, and must be able to guarantee performance and reliability to their customers. In converged voice and data networks, the network infrastructure must deliver very high quality and availability for some customer needs, while also providing low-cost high-capacity bandwidth for other needs. The use of quality of service mechanisms to provide prioritization for various traffic types is a key element needed for voice and data network convergence. However, it is not sufficient if the underlying networks are unreliable. The focus of this article is to address the reliability aspects of VoIP services, including the underlying IP networks.

## INTRODUCTION

Converged voice and data networks need to be able to meet varying service needs for different market segments. Many users have come to expect highly available telecommunications services supporting high-quality voice and accurate data. In converged IP networks, the network infrastructure must deliver very high quality and availability to meet these customer needs.

While IP technology provides the means to offer multiple grades of voice quality, the focus of this article is to address the reliability aspects of high-end voice over IP (VoIP) services, including the underlying IP networks. To provide a basis for VoIP reliability, an overview of a VoIP functional architecture is provided. VoIP service reliability is described through downtime and defects per million metrics. The key factors that impact the reliability of IP voice services and networks are covered. The issues are described in terms of the elements needed to provide VoIP services on a converged network. The current state of IP network reliability and the current challenges along with potential solutions for improving IP reliability are described. Methods of designing reliable VoIP networks

are described, along with restoration mechanisms. Finally, current trends in assessing end-to-end VoIP reliability are presented.

## A VOIP SERVICES ARCHITECTURE

Figure 1 provides a high-level view of a VoIP functional architecture. This is a logical architecture and is not intended to represent the physical implementation. Depending on the particular network implementation, some of these network components may be combined or separated (e.g., a combined signaling and trunking gateway).

## VOIP VALUE-ADDED SERVICES

Beyond traditional voice services, VoIP has encouraged the emergence of new applications that maximize the simultaneous use of voice and data communications. A wide range of these applications is now available, and new applications are emerging [2]. Some examples include unified messaging, distributed call centers, custom call handling, IP Centrex, IP private branch exchange (PBX), and voice virtual private network (VPN).

## QUALITY OF SERVICE FOR VOIP

One of the key requirements for the widespread deployment of VoIP is the ability to offer toll-quality service equivalent to the existing public switched telephone network (PSTN). The quality of a telephone call is primarily a function of distortion and delay. Distortion is the difference between the received and original signals. Delay is the time elapsed from the origination of the speech signal until the destination user receives it. While the PSTN network ensures fixed-delay minimum-distortion service, this is not necessarily the case for IP-based networks. IP networks have evolved around best effort service and typically do not provide guarantees for key performance criteria. The need to support real-time services has driven the development of control mechanisms and technologies that can be chosen to provide quality of service (QoS) support similar to constant bit rate (CBR) and real-time variable bit rate (VBR-rt) in asynchronous transfer mode (ATM). The objective is always to guarantee prioritization of voice flows over best effort data, and ensure that their performance is not compromised by traffic congestion.

In order to make the transition to high-end profitable VoIP services, service providers need reliable high-performance networks and must be able to guarantee this reliability and performance to their customers. High-end VoIP services will only be successful if they are as reliable as the alternatives they replace. On the other hand, IP technology also provides the capability to offer lower-grade, and more economical, voice service. Service differentiation in the form of service classes backed up with strict service level agreements (SLAs) requires the underlying network to be reliable. In the next section we focus on the reliability aspects of VoIP.

## VOIP SERVICE RELIABILITY

PSTN callers expect to be able to succeed in making a call every time they pick up the phone. They also take for granted high voice quality in typical phone conversations. What will it take for VoIP to finally live up to its promise of radically simplifying telecom networks? The answer comes down to one word: reliability. When VoIP networks achieve PSTN-like availability, carriers will benefit from simplification of their network architectures, simplification that will lead to improved profitability.

This section addresses the emerging customer-oriented definition of VoIP service availability. It also describes how the defect per million (DPM) approach can be used to quantify the users' end-to-end VoIP perspective.

### CARRIER CLASS AND "FIVE-NINES" RELIABILITY

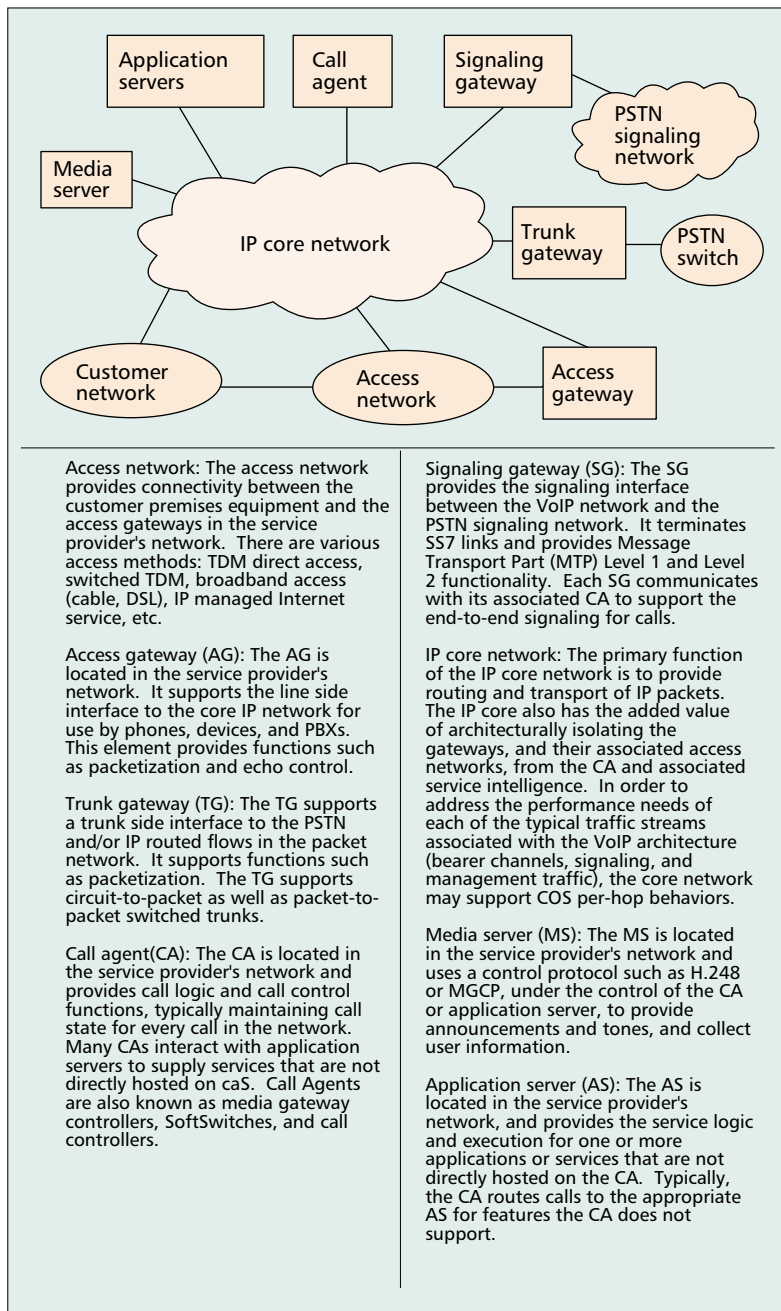
The term *carrier class reliability*, which comes from the PSTN world, is most often associated with 99.999 percent availability, which amounts to about 5 min downtime/year. When applying the concept of carrier class availability to VoIP networks, it is important to take a closer look at how this is defined and also at the design of the network.

In the PSTN, different elements have differing availability objectives based on the function they perform, element redundancy, and overall service architecture. For example, class 4/5 switches require 99.9996 percent availability and are typically single points of failure. In contrast, databases may only need 99.99 percent availability if they are replicated in the network with fast failover mechanisms. The goal for VoIP is to achieve the same level of network reliability as provided by PSTN infrastructure.

### THE CONCEPT OF RELIABILITY FOR VOIP SERVICES

Availability is well defined for component and system design, as described in many standard texts. There is an underlying concept of up-state, during which the system can perform its function, and down-state (outage), during which the system cannot perform its function. However, for VoIP services, the concept of availability is not as easily applied.

For VoIP service, the system is the entire network (or set of networks) and associated systems designed to provide the service between two end users. The function is a combination of connection setup, transfer of user information,

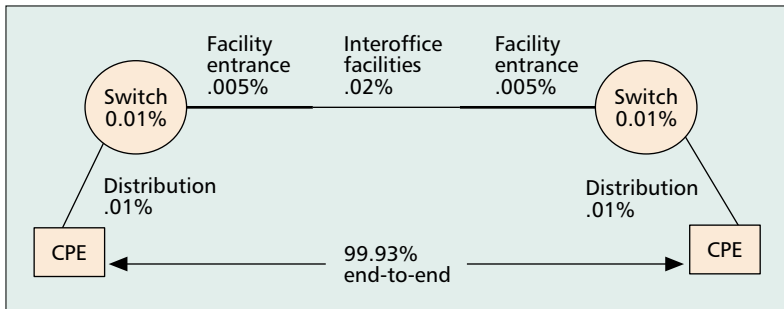


■ Figure 1. VoIP functional architecture [1].

and maintenance of the connection. Only if the logical connection can be completed with a reasonable reattempt algorithm, the QoS is sufficient for useful communications, and the calls can be maintained for sufficient time to complete the transaction, is the communication service defined to be in up-state and available. The fraction of time during which the service is available between a pair of end points is called service availability. Reliability metrics for VoIP services are described in the next section.

### RELIABILITY METRICS FOR VOIP

First, key metrics to quantify the reliability of VoIP services are defined. Then objectives can be established based on these metrics to ensure high reliability. The telecommunications commu-



■ **Figure 2.** *PSTN reference connection.*

<b>Accessibility</b>	Ability to initiate a voice call when desired
<b>Continuity</b>	Successful continuation of a call to its completion with no interruption, given that accessibility has been successful
<b>Fulfillment</b>	Ability to deliver call service that meets the user's quality expectations, given that the call completes successfully

■ **Table 1.** *VoIP service defects.*

nity introduced a QoS Framework (QSF) to define the quality and reliability of PSTN-based services [3]. The QSF consists of a matrix with rows made up of communication functions performed by service users, and columns made up of quality and reliability criteria perceived by users. Communication functions were classified into three major categories: *connection establishment*, *user information transfer*, and *connection release*. Quality and reliability criteria included *speed*, *accuracy*, and *reliability*.

A new framework can be developed to define, measure, and monitor end users' experience of quality and reliability of VoIP services. Users need to be able to access a voice application, initiate a voice call, continue using the voice application with no interruption for a desired duration of time, and fulfill the initiated voice call transaction at an acceptable quality. For VoIP service reliability, we combine all these criteria into two reliability metrics: *end-to-end downtime* and *DPM*.

**End-to-End VoIP Service Downtime** — Current PSTN customers expect a high level of network reliability based on their past experience. To maintain this high level of reliability, the VoIP service reliability objectives should preserve these expected levels. The end user will not care, and in many cases will not know, what network technology and architecture is employed.

An end-to-end availability model for VoIP may be derived from the classic PSTN end-to-end availability model [4]. The simple high-level architecture for the PSTN is shown in Fig. 2. It models a *single path* through the PSTN from one end user to another, including the major network segments needed to connect two telephone subscribers served by two different central office switches. In this model, the availability requirement for the originating and terminating loop was 99.99 percent, equivalent to annual downtime of 0.01 percent, or about 53 min.

Each segment of the telephone network was

engineered to meet specific availability requirements, yielding end-to-end availability of 99.93 percent, or about 365 min/year of downtime for a single path. For a particular VoIP network implementation, an end-to-end VoIP reference connection and an associated reliability model can be developed. To make VoIP reliability comparable to the PSTN requires the availability of VoIP elements to be much higher than 99.93 percent. The network implementation and elements availability will determine end-to-end availability. Higher end-to-end availability can be achieved by element redundancy and network diversity.

**Defects Per Million** — End-to-end downtime is the traditional customer-oriented measure of unavailability. However, this measure does not reflect the impact of outages on customers because it does not incorporate customer demand during outages. Defects per million is a reliability metric that only counts customer demands not served. The average number of blocked calls and cutoff calls per million attempted calls is referred to as DPM. A blocked call is defined as a call that was prevented from being successfully set up or completed due to failures. A cutoff call occurs when a stable call is terminated prior to either party going "on-hook." The fact that large carriers transport several hundred million calls per day implies averaging over a large sample in DPM calculation.

There are three sources of VoIP service defects, summarized in Table 1.

The total network DPM is approximated by the sum of DPM in accessibility, continuity, and fulfillment. It should be noted that voice quality monitoring is essential to capture VoIP service fulfillment DPM. It is important to measure the voice quality users experience to assess call defects.

For VoIP service fulfillment, performance thresholds must be specified to define what constitutes a defective call. These thresholds are set based on user tolerance to degraded voice quality as well as technology limitations. The parameters that contribute to call impairments are random packet loss, consecutive packet loss, excessive delay, and delay variation. The thresholds for these parameters are functions of many other parameters. For example, the impact of packet loss in voice quality is dependent on the particular codec in use and, more significantly, highly dependent on the packet loss concealment algorithm.

## QOS TECHNOLOGIES TO SUPPORT VOIP RELIABILITY

Once the definition of service reliability extends beyond the simple question of whether the service up or down to include acceptable performance, one has to consider questions like "Is there enough capacity to support the expected load on this route?" and "Is there enough bandwidth allocated to voice service?" This leads to a class of issues and methods related to guaranteeing QoS to a specific service on a shared network. Since this is not the main topic of this

Method	Goals/characteristics	Status in standards
Constraint-based routing and traffic engineering	Aim at selecting the paths on which voice calls are routed to guarantee acceptable delays as well as reserving sufficient capacity for the expected load.	The leading technologies for performing this function are MPLS with Traffic Engineering (MPLS-TE) [5] combined with extensions to the IP routing protocols, like OSPF-TE and ISIS-TE. MPLS-TE can also be used for configuring restoration paths for link and router failures.
Bandwidth management and admission control	Handle requests for bandwidth allocation for VoIP calls and limit the number of calls in progress to comply with the allocation.	The leading current technical direction is to use the RSVP-TE [6] protocol to set up paths with reserved bandwidth. There is current work in the IETF on DS-TE models and algorithms that will implement the call admission function [7].
QoS for established voice calls	Ensure that accepted calls are given the proper priority so that the committed QoS level can be met even if a link is congested due to non-real-time traffic.	This function is provided by IP DiffServ [8] mechanisms, like packet differentiation and per-hop behavior, which provide priority queuing and packet dropping rules at each node along the path, including a specific priority queue for real-time traffic.
Adaptive coding	Reduce bandwidth consumption under congestion and still provide good QoS to calls in progress.	Requires enhancements to SIP.


■ **Table 2.** *Methods for providing VoIP QoS.*

article, Table 2 provides only a high-level description of the leading methods and mechanisms, some of which are still under discussion in the Internet Engineering Task Force (IETF) and the technical community at large.

## IP NETWORK DESIGN FOR RELIABILITY

IP networks were initially designed to carry best-effort data traffic, which can tolerate short disruptions (up to several minutes) and increased delay and packet loss caused by rerouting at the IP layer around failed facilities or routers. These relatively relaxed requirements on outage time had a profound impact on the design of IP networks and routers. This section presents the trade-offs in IP network design between node reliability, path diversity and redundancy, and restoration times, and how they all impact the overall availability of IP networks.

Due to the inherent lack of varrier grade reliability for IP backbone routers, network designers have to ensure dual physically diverse paths when constructing the topology of large-scale IP networks. Adequate capacity over alternate routes needs to be provisioned to allow for rerouting packets under all types of failure conditions (e.g., backbone router failure, transport fiber cuts). Backbone routers are typically collocated in pairs, linked to each other. Network designs involve trade-offs between desired reliability and cost of capital expenditures. A fully redundant design involves each of two collocated backbone routers in an office linked to one backbone router in an adjacent office with sufficient capacity to route all traffic flows in the event that one backbone router fails (Fig. 3). Such designs guarantee restorability for all traffic flows and do not introduce additional delay for the rerouted traffic, but involve significant expense. For that reason, capacity designs optimize cost while ensuring sufficient spare capacity



■ **Figure 3.** *An example of a fully redundant backbone.*

for the desired level of restorability; however, the rerouted path may involve additional path delays and packet loss.

Currently, restoration of packet flows over alternate routes under failure conditions follow IP distributed routing protocols such as Open Shortest Path First (OSPF). Reroute times typically average tens of seconds and are not adequate for VoIP services. As newer packet protocols such as multiprotocol label switching (MPLS) are introduced, efforts to derive faster restoration schemes using MPLS traffic engineering are under way. For example, MPLS fast reroute schemes identify backup paths in advance for any given path, such that in the event of failure, the backup path is immediately invoked. Initial claims on such schemes suggest that subsecond restoration is possible. Issues such as the practical implementation of such schemes in large networks are still under investigation. For example, it may not be possible from

*Without redundancy, hardware component failures interrupt customer traffic until the failed component is recovered by reset, which typically takes several minutes, or until it is replaced, which can take hours.*

a cost or reliability perspective to apply fast reroute schemes to all classes of traffic in a large network. That would lead to some form of traffic segmentation; VoIP, for example, would qualify for fast reroute, whereas best effort IP traffic would continue to be restored over existing OSPF mechanisms.

Finally, the current role of transport-related restoration mechanisms (synchronous optical network [SONET] rings/optical wavelength switching) appears to be limited in IP networks, even though such restorations are already in the subsecond range. Current backbone routers are not considered carrier grade in terms of reliability. The failure of a backbone router cannot be recognized at the physical transport layer today. Thus, an efficient spare capacity design for all backbone router failures can also result in ensuring spare capacity for the vast majority of link/fiber span failures. If, however, future backbone routers can be designed as carrier grade in terms of reliability, only link/fiber span failures need be considered. In such conditions, transport-based restorations become feasible from a cost perspective. Also, future protocol advancement under the generalized MPLS (G-MPLS) umbrella may be able to blur the distinction between packet and optical layers so that more efficient (and fast) restorations can be considered.

## VOIP ELEMENT RELIABILITY

### IP ROUTER RELIABILITY

From a reliability analysis perspective, the IP network consists of routers interconnected by facilities. This section provides a brief overview of the main router elements with their typical mean times between failures (MTBF) and mean time to restoration (MTTR) as well as redundancy mechanisms, which are important for availability modeling of services in IP networks [9]. A high-speed IP router is a special multi-processor system with two types of processors: route processor (RP) and line card (LC). The RP controls the operation of the entire router, and the LC interfaces facilities, or links connected to other routers, servers, or VoIP elements. A switch fabric interconnects route processors and line cards. These hardware (HW) components share a chassis and operate under the control of an operating system (OS) representing the software component of the router. A typical MTBF for RP and LC is in the range 60,000–150,000 h. This MTBF accounts only for hard failures, which require replacement of the failed component, in contrast with soft failures, from which the router can completely recover (e.g., by card reset). A typical example of a soft hardware failure is parity error. Router vendors do not usually provide software MTBF, but existing experience suggests that it is below 100,000 h (sometimes by an order of magnitude). Software MTBF does not account for soft HW failures. They are counted only in MTBF for all failure types.

Without redundancy, HW component failures interrupt customer traffic until the failed component is recovered by reset, which typically takes several minutes, or replaced, which can

take hours. To reduce failure impacts, shared HW components that impact the entire router if failed (e.g., RP, switch fabric, uplink LC) are typically redundant. Then the restoration time (assuming a successful switchover to the redundant component) is defined by the switchover time. All router vendors are pursuing high-availability programs with the goal of implementing hitless switchover to the redundant component. Hitless switchover refers to a switchover done within a few seconds (typically below 5 s) in contrast with a usual switchover, which takes a few minutes. Switch fabric, RP, and uplink LC have different redundancy and restoration mechanisms. Switch fabric has  $n$  active elements, and failure of one element triggers load reallocation to the remaining ( $n - 1$ ) elements. Restoration time is on the order of 1 s. RP redundancy is provided by a configuration with two RP cards: primary and secondary. Hitless restoration from RP failures requires synchronization of routing tables between primary and secondary RPs. There are two different synchronization solutions known as stateful switchover/non-stop forwarding (SSO/NSF) and non-stop routing (NSR). SSO/NSF requires NSF awareness on all neighboring routers, while NSR provides restoration without communication to the neighboring routers. Uplink LC redundancy is provided by two uplink cards connected by diverse links to two backbone routers. Traffic is evenly distributed between the two uplinks. Restoration from uplink LC (or uplink itself) failure occurs using an OSPF restoration mechanism, and the restoration time is on the order of 10 s. Successful restoration without traffic loss requires a capacity management mechanism for keeping single uplink utilization below 50 percent. For customer facing cards, LC redundancy is currently available only for SONET interfaces, and usually only 1:1 redundancy is supported with switchover time to the redundant card on the order of 1 min. With all of the above redundancy in place, the downtime contribution from RP and LC failures (both HW and software) becomes less than 5 min/yr, which corresponds to better than five nines of availability. However, the current router availability is between three and four nines because of the following factors:

- The chassis remains a single point of failure. Even with MTBF in the range 40–50 yr, chassis contribution to downtime is on the order of 10 min due to large MTTR.
- Abbreviated development and testing cycles have led to a trade-off between reliability and time to market. Many defects are discovered only after deployment of a new card or software release. Some of these defects cause failures that cannot be mitigated by redundancy.
- Several software and HW upgrades per year are typical, with each upgrade contributing 10–60 min of downtime.
- Distributed denial of service attacks (DDOS) may cause long outages.

To achieve five nines of router availability, significant progress must be made in improving router design, quality of HW and software,

developing hitless software and HW upgrades, and better procedures for preventing and detecting DDOS.

### RELIABILITY TECHNIQUES FOR VOIP

A VoIP network distributes the functionality of the classical central office (CO) switch of the PSTN throughout the network. The stringent PSTN product reliability requirements assume specific network and product architecture and technologies. PSTN design strategy is based on highly fault-tolerant high-capacity circuit switching architecture. However, VoIP network elements vary in size and function, use different protocols that are more tolerant to failure, and are deployed in a network with redundancy. Consequently, PSTN element reliability requirements should not be used for all VoIP elements. Reliability requirements for VoIP elements should be determined based on failure mode impacts. VoIP element reliability requirements can then be determined based on element size and network design strategy.

There are multiple techniques available to meet high-availability VoIP service requirements. These techniques include fault-tolerant HW, fault-tolerant software, system redundancy, and network interface redundancy.

**Fault-tolerant HW** is the traditional approach to achieve high availability in standalone systems. With this technique, redundant HW is built into the platform, and upon failure switchover to the redundant hardware occurs seamlessly.

**Fault-tolerant software** relies on monitoring of software elements, and the affected functions are transferred to another process upon detection of a problem.

**VoIP element redundancy** replaces the high reliability of individual elements with having enough elements to ensure high service availability. In order to detect failures, messages commonly known as *heartbeats* are exchanged either between elements or with an operations support system.

**VoIP element interface redundancy** employs multiple network interfaces and the ability to switch between the interfaces in real time.

In summary, the effective deployment of the VoIP elements must account for a variety of potential failures. The likelihood of different failures should be considered when evaluating various architecture options:

- High availability requires multiple but diverse network paths between all elements. In the event of a network failure, a fast restoration capability (< 2 s) is required so that calls in progress are not terminated.
- Software failures are overcome with redundant systems that are able to execute different software versions, yet interoperate and failover in a matter of seconds.
- HW failures are also overcome with redundant elements in a cost-effective manner. While fault-tolerant HW does offer an additional level of protection, its cost-benefit equation is still debatable. The additional cost may not be deemed worthwhile since the required functionality can be delivered via other mechanisms.

## VOIP RELIABILITY MONITORING AND ASSESSMENT

Assessing end-to-end VoIP service reliability depends on the ability to effectively monitor and correctly assess the reliability of the transport network and service elements. Current methods of fault isolation and root cause analysis still rely on manual effort and are far from fully automated. The IP networks of today generate vast numbers of alarm messages for a variety of conditions (including element outage). These messages tend to be brief and cryptic; hence, significant and time-consuming human evaluation is necessary to identify the onset of an element outage and understand the reliability impact. It is therefore imperative to improve the quality, precision, and ease with which network element outages are reported and acted on by service providers.

Efforts are currently underway in standards bodies to standardize the automated delivery of physical element outage reports in packet networks. Complete automated reports detailing the identity of the failed element and the outage duration time can significantly enhance the monitoring of an element's life cycle reliability.

Efforts in standards bodies have commenced to provide accepted methodologies to estimate network availability. Agreement has been reached on the estimation of access availability in IP networks based on the state of customer facing ports (available or unavailable) over a defined period of time [10]. Access availability can then be defined as the ratio of the total available time for all customer ports in the network to the total port time in a defined period (e.g., one month, one year). A time-based access network DPM, which is a scaled version of availability, is derived with the appropriate scaling factor. Port availability can also be further weighted by their bandwidth or any other utility function. Backbone availability, the ability of packet flows to successfully reach their destination upon successful network access, is more complicated. The complication arises from the multiple paths available to packet flows between any pair of access and egress routers. Current efforts involve an averaging of a reference path through the backbone in an IP as well as MPLS environment.

An improved outage reporting mechanism and better understanding of network availability can go a long way toward the delivery of the desired end-to-end VoIP service availability. Efforts are underway in standards bodies to reach agreement on the definitions of acceptable performance/quality for VoIP services. The goal is to drive vendors to develop standardized measurement mechanisms for acceptable voice quality.

## CONCLUSION

Providing high reliability for VoIP services is a challenging and multifaceted problem. While some customers may be willing to trade off performance and reliability for lower-cost solutions, many will not. Therefore, VoIP networks must be able to meet the needs of both types of users.

*Efforts are underway in standards bodies to reach agreement on the definitions of acceptable performance/quality for VoIP services. The goal is to drive vendors to develop standardized measurement mechanisms for acceptable voice quality.*

Metrics and monitoring capabilities are essential to manage VoIP reliability. Two key metrics, downtime and DPM, are presented for tracking element and service reliability. Industry agreed standards should drive vendor equipment improvements.

There are many reliability issues to be resolved in order to achieve acceptable reliability for high-end VoIP services.

Reliability of the underlying IP networks is a key element. As discussed in this article, IP router reliability must improve through better designs, higher-quality hardware and software, and hitless upgrades. In addition, better procedures for preventing and detecting DDOS and effective QoS mechanisms are needed to reduce impacts on VoIP services.

VoIP network element reliability is also critical in providing highly reliable services. While many products are in the early stages of their life cycle, design techniques from more mature technologies are applicable. Failure detection and recovery mechanisms at the element and network levels are a necessity. Metrics and monitoring capabilities are essential to manage VoIP reliability. Two key metrics, downtime and DPM, are presented for tracking element and service reliability. Industry agreed standards should drive vendor equipment improvements.

Robust network designs coupled with fast restoration mechanisms are critical to the success of VoIP technologies and services.

## REFERENCES

- [1] GR-929-CORE, "Reliability and Quality Measurements for Telecommunications Systems," issue 8, Dec. 2002.
- [2] W. J. Bushnell, "VoIP Business Services — What and How," *Business Commun. Rev.*, Mar. 2001.
- [3] J. S. Richters *et al.*, "A Framework for Defining the Quality of Communication Services," *IEEE Commun. Mag.*, Oct. 1988.
- [4] C. M. Hamilton *et al.*, "Telecommunication-Network Dependability: A Baseline on Local-Exchange Network Availability," *IEEE Reliability and Maintainability Symp.*, 1991.
- [5] D. Awduche *et al.*, "Requirements for Traffic Engineering over MPLS," RFC2702, Sept. 1999.
- [6] D. Awduche *et al.*, "RSVP-TE: Extensions to RSVP for LSP Tunnels," RFC 3209, Dec. 2001.
- [7] F. LeFaucheur and W. Lai, "Requirements for Support of DiffServ-Aware MPLS Traffic Engineering," RFC 3564, July 2003.
- [8] S. Blake *et al.*, "An Architecture for Differentiated Services," RFC2475, Dec. 1998; also D. Grossman, "New Terminology and Clarifications for DiffServ," RFC 3260, Apr. 2002.
- [9] M. Vogt, R. Martens, and T. Andvaag, "Availability Modeling of Services in IP Networks," *Design of Reliable Commun. Networks 2003*, Banff, Alberta, Canada, Oct. 19–22, 2003, pp. 167–72.

- [10] "Access Availability of Routers in IP-Based Networks," Committee T1 tech. rep. T1.TR.78-2003.

## BIOGRAPHIES

CAROLYN JOHNSON (crjohnson1@att.com) is a technical manager in Network Design and Performance Analysis at AT&T Labs, Middletown, New Jersey. She has extensive experience in telecommunications performance and reliability analysis. She has modeled voice services performance, evaluated the introduction of new network technologies, and designed congestion control algorithms. Her recent work is focused on VoIP technologies and services. She joined Bell Laboratories after receiving her Ph.D. in mathematics from the University of Florida in 1980.

YAKOV KOGAN [F'01] (yaakovkogan@att.com) is a technology consultant at AT&T Labs, Middletown, New Jersey, in the Network Design and Performance Analysis Department. After receiving a Ph.D., he worked on performance analysis of computer and communication systems, and developed nonparametric and asymptotic methods for solving stochastic models of large dimension. His recent activities include performance and reliability analysis of corporate IP backbone networks. He is a member of IFIP Working Group 7.3 on Computer System Modeling.

YONATAN LEVY (ylevy@att.com) is a technology leader in the Network Design and Performance Analysis Division of AT&T Labs. He has vast experience in performance modeling and analysis, holds several patents on dynamic network call distribution and delivering QoS in packet networks, and has more than 20 publications. In September 2000 he organized an ITC specialist seminar dedicated to IP traffic. He received an M.S. in operations research in 1973 from Tel-Aviv University and a Ph.D. in mathematical sciences from The Johns Hopkins University in 1980.

FARHAD SAHEBAN (saheban@att.com) is a principal technical staff member at AT&T Labs. He joined AT&T Bell Laboratories in 1986 after he graduated from the University of Southern California, Los Angeles with a Ph.D. in electrical engineering. He has over 25 years of industrial and academic experience, including design of fault-tolerant systems and networks, self-repairable multiprocessor systems, and built-in self-test systems. His current areas of focus include VoIP reliability and QoS, network reliability, and software reliability engineering.

PERCY TARAPORE (tarapore@att.com) is a technology consultant in Telecommunications Standards at AT&T Labs. He has extensive experience in reliability analysis and performance modeling of network services. He developed models for estimating blocked calls from network failures and optimized capacity for AT&T's FASTAR restoration network. Currently he represents AT&T in reliability and performance related standards for IP applications. He received his M.Sc. degree in physics from Bombay University in 1976, and an M.S. degree in systems engineering from The Ohio State University in 1981.