

Best Practices for Planning, Deploying and Managing Enterprise VoIP



April 2008

Introduction

Voice over IP (VoIP) technology has matured rapidly since its inception roughly 10 years ago. Based on this maturity, VoIP is now enjoying an accelerated acceptance by the market, with the majority of IT organizations deploying VoIP or planning to do so in the near future¹. There are many factors driving the adoption of VoIP. In addition to the potential cost savings, VoIP offers the prospect of facilitating the integration of voice communications with strategic data applications to increase worker productivity and gain competitive advantage.

Early adopters have learned that real-time applications, such as VoIP, are far more demanding of predictable network performance than the typical enterprise data application. Successful VoIP deployment often requires fairly substantial modifications of the network, plus careful monitoring of all traffic to verify quality voice service and adequate performance for data applications. However, most enterprises are undaunted by the challenges of optimizing their networks for VoIP in large part because they have already embarked on optimizing application delivery for their rapidly expanding suites of mission-critical enterprise data applications and services that are becoming much more closely intertwined with fundamental business processes.

The goal of this white paper is to help the reader gain a better understanding of the challenges presented by VoIP as well as what must be done from a planning and management perspective to successfully deploy VoIP. The discussion includes a high level description of a proven methodology for deploying VoIP and other demanding applications, as well as a discussion of a management system that is optimized for both fault and performance management of converged voice and data networks.

The Migratory Movement to Implement VoIP

VoIP technology has reached a point of maturity where it is now both the preferred and inevitable solution for new deployments of enterprise telephony. In spite of this dramatic movement away from deploying traditional TDM-based PBXs, most organizations have opted to preserve their investments in TDM-based PBXs and phones via a migration to VoIP by installing hybrid technology that supports a mix of legacy TDM and VoIP lines. According to Infonetics², hybrid VoIP/TDM PBXs are expected to hold an estimated 63% share of the PBX market in 2007 vs. approximately 18% for pure IP devices and another 19% for TDM products.

Early adoption of VoIP was driven primarily by cost savings in several dimensions:

- Reduced long distance toll charges
- Reduced cost of moves, adds and changes
- Lowered cost of IP PBXs vs. TDM PBXs

- Reduced number of PBXs through better scalability of IP PBXs
- Elimination of parallel data and voice wiring systems

As the technology continues to mature, VoIP is now viewed more strategically because of its ability to improve business operations and worker productivity by enabling mobility, presence, unified messaging, and video conferencing. According to the 2005/2006 VoIP State of the Market Report from Distributed Networking Associates³, strategic voice and data applications now constitute the number one reason for adopting an IP PBX solution, followed closely by the cost savings described above.

Of the enterprises that are deferring adoption of VoIP, many are generally concerned that the initial capital expenses (CapEx) associated with the transition will outweigh the savings in operating expenses (OpEx) and the strategic advantages that their particular organizations would gain. Another common impediment to VoIP deployment is the perception that there is a lack of adequate solutions for managing and troubleshooting VoIP quality. For example, the 2005/2006 VoIP State of the Market Report documented that the lack of systems for managing and troubleshooting VoIP quality was tied with security concerns as the number one impediment to deploying VoIP. In addition, the primary reason that enterprises consider a hosted VoIP service from a Managed Service Provider (MSP) is to avoid personal responsibility for the management and administration of the IP telephony system.

In spite of concerns about VoIP management, it is worth noting that IT organizations are actually reluctant to deploy new, special-purpose management tools. For this reason, enterprises often make the mistake of deploying tools for tracking and managing VoIP performance only after a significant VoIP implementation has occurred and performance issues have been encountered.

The Challenges of Deploying VoIP

As IT organizations embark on the implementation of VoIP across their LANs and WANS, they face of a number of challenges:

Quality of Voice Services

End users expect VoIP audio quality to be equal to that delivered by circuit-switched PBXs and public switched telephony networks (PSTNs). Toll quality voice is generally equated to a subjective Mean Opinion Score (MOS) of 4 on a 5 point scale. In legacy TDM and circuit-switched networks, voice quality may be impaired by problems such as echo, distortion from digitization, voice level variation, and background noise. But these issues have been addressed and practically eliminated through 100 years of R&D. In the new, packet-switched VoIP network, these problems have not only reappeared, but further impairment may be caused by unpredictable delay due to packet processing and buffering, variation in packet delay (also known as jitter), and the loss of packets.

Packet Delay and Delay Variation

In the legacy telephony system, the one-way, end-to-end terrestrial delay is a constant sum of codec delay, switching delay, and propagation delay. As a result, there is little difficulty in keeping delay under 150 ms, which is the maximum delay normally associated with a MOS of 4 or better. With VoIP delay is compounded by packet processing and serialization, packet buffering when links are congested, and intentional buffering to remove delay variation or jitter. In a packet-switched network, consecutive packets can experience widely different delays due to a number of reasons, including variable queuing delay due to congestion, variable packet switching delay, and abrupt changes in routing topology. Jitter becomes a serious problem when it approaches 75 ms. For this reason, most VoIP receivers (i.e. phones and gateways) include a de-jitter buffer that generally delay packets by an additional 40 ms to allow a steady stream of packets to be delivered to the codec. The 40 ms de-jitter buffer delay must be factored into the 150 ms delay budget.

Packet Loss

The occasional loss of individual packets or strings of packets is a result of IP bit errors and buffer overflows during periods of congestion. In addition, if packets arrive too late at the de-jitter buffer, they may be dropped. Because of the real-time nature of voice and the limitations on allowable delay, it is not possible to retransmit dropped packets. While VoIP quality is not very sensitive to loss of isolated packets, it is quite sensitive to losses of strings of consecutive packets lasting for more than 200 ms. As a result, it is somewhat difficult to equate long term averages of packet loss rates to voice quality. A low average loss rate may obscure the fact that strings of packets are being dropped. On the other hand, average loss rates as high as two or three percent might be tolerable where loss is restricted primarily to isolated packets. For isolated packet loss, some VoIP codecs include a feature called packet loss concealment (PLC) that reconstructs missing packets within a VoIP packet stream. PLC uses a predictive speech model to synthesize the payload of missing packets, compensating for gaps of up to 20 to 40 ms with no discernable effect on the actual quality or the subjective MOS rating.

In order to ensure that the VoIP impairments (delay, jitter, and packet loss) remain within acceptable bounds it is necessary to implement Quality of Service (QoS) functionality in a consistent fashion across all packet forwarding devices (switches and routers) in the network. Ideally, QoS functionality can be configured to ensure that voice traffic receives an appropriate share of total bandwidth and that voice packets are priority-queued whenever congestion occurs. Priority-queuing means that voice packets in the queue are always sent before any non-voice packets. Once QoS has been configured, careful monitoring of VoIP call quality is needed to ensure the effectiveness of the chosen QoS policy and its implementation.

Call Setup and Signaling

With circuit switching, signaling is used to create an end-to-end, dedicated communications channel that remains in place for the duration of the call. While Signaling System 7 (SS7) is the signaling method used in the PSTN, PBXs often use a variety of signaling methods for call setup.

In the case of VoIP, call signaling is typically implemented using the H.323, SCCP or SIP protocol. These call signaling protocols are typically used to determine the admissibility of a call being placed or to determine the presence of the called party. Signaling is also used by the endpoints to negotiate a common set of call parameters, such as the codec algorithm to be used and the source and destination port numbers that will be used for the duration of the conversation. Since the VoIP call set up process involves the IP PBX as well as the endpoints, call setup times may exceed user expectations during periods of heavy call volumes or network congestion. As a result, QoS policies need to be enforced to provide high priority for both call signaling traffic and call bearer traffic.

Service Availability

In legacy voice networks, PBXs and trunk circuits are generally designed to provide service availability in the 99.99% to 99.999% range. Users typically expect a similar level of availability for VoIP in spite of the fact that few enterprise LANs or WANs are designed to meet this high standard of reliability and resiliency. IP networks that provide availability in the 99.99% to 99.999% range are built in part by minimizing single points of failure and using redundant, highly resilient network devices. However, improving the availability of an IP network requires more than just a robust network design. Improving availability also requires IT organizations to significantly improve their approach to key management tasks such as performance, service, fault and configuration management.

Coexistence with Data Applications

QoS can be viewed as a zero sum game because minimizing delay, jitter, and packet loss for voice traffic can come at the cost of reducing the service level for data applications that share the network. In the LAN, the effects of VoIP prioritization on data application performance are generally negligible. Over narrow band WAN links, however, the percentage of voice traffic needs to be controlled in order to prevent starvation of mission-critical data applications. Properly implemented, QoS policy establishes traffic classes for real-time applications like voice, yet protects mission critical data applications from being overly impacted by the real-time applications. Best practices for QoS implementation usually involve limiting VoIP traffic to less than 30% of normal link utilization. Applying such rules of thumb, however, is not a substitute for careful monitoring and testing the performance of all mission critical applications.

Network Readiness

The basic requirements of an IP network that can accommodate VoIP and other real-time applications include:

- Effective network management tools and processes
- LANs based entirely on full duplex Layer 2 and Layer 3 switching

- Adequate bandwidth provisioned on all LAN and WAN links
- End-to-end QoS implementation over the LAN and WAN
- High availability network design based on redundancy and resiliency

In many cases, successful VoIP deployment will require that IT organizations modify their approach to network management. In addition, some fairly significant modification of the underlying IP network infrastructure will be required as a prerequisite for deployment of VoIP.

Converged Organizations

As a result of VoIP deployment, traditionally separate groups within IT find themselves working side-by-side. Telecom staff, accustomed to dealing with the traditional circuit-switched PBX and phone company issues, must work closely with network engineers whose expertise is focused on IP and packet-switched routing. How well this organizational convergence plays out depends not only on managerial commitment to cooperation, but also on adopting network management tools that integrate the metrics of interest to each team to ensure adequate performance for all applications and network services. Converged management tools for the converged network often prevents finger-pointing when problems develop and resolution procedures don't reveal an obvious cause.

A VoIP Management Methodology

In order to meet the challenges of deploying VoIP, organizations responsible for converged data networking and telecommunications need to adopt a comprehensive methodology for planning, managing, and operating the network as an application delivery platform that meets user expectations for availability and performance of all applications, including VoIP.

A general approach to the management of application delivery is described in the 2008 Application Delivery Handbook⁴, published by Ashton, Metzler and Associates. This document stresses the fact that successful delivery of both data applications and real-time applications requires a life cycle project and network management model that involves the following activities:

- Planning
- Deployment and Application Optimization
- Management
- Control

Planning

If VoIP deployment is not adequately planned, serious disruptions to voice and data communications are possible, together with a prolonged period of transition from the legacy system to VoIP. Planning for VoIP involves the following activities:

VoIP Characterization

Working with the chosen VoIP vendor(s), the planning team should characterize the VoIP traffic expected to flow over the LAN and WAN. The characterization should include setting target thresholds for delay, jitter and packet loss, as well as establishing estimates of the traffic loading on the LAN and WAN. Impairment thresholds will be somewhat dependent on the functionality and configuration of VoIP endpoints, including codecs, de-jitter buffering design, QoS classification setting(s), and PLC capability. Establishing a target for inter-site and intra-site availability and latency is another important aspect of this initial planning.

Network Assessment

As outlined in the previous section of this document, a careful analysis of the readiness of the existing network for voice communications is required. The VoIP characterization and the chosen model for QoS implementation should be used to provide detailed guidelines for this analysis. In most cases, some degree of network redesign and modification will be required to ensure the desired level of availability and performance. An adequate network assessment involves a good understanding of the baseline performance and capacity of the existing network, as well as some form of network modeling for estimating the performance improvement expected from planned enhancements or additional capacity to the network.

VoIP Impact Analysis

VoIP deployment can potentially impact mission critical interactive applications, especially in the WAN where voice and data traffic may need to share narrow band links. Modeling the impact of VoIP traffic on the performance of mission critical applications may require further adjustments to the network design and/or the QoS implementation.

Management Readiness

Ensuring that management processes will be ready for VoIP deployment is another aspect of the planning process. In addition to addressing the organization issues of data communications vis à vis telecommunications, an assessment of management readiness should identify any requirement for additional staff members, technical training, or management tools that are compatible with the chosen VoIP solution(s). Planning for the management of VoIP may also include consideration of outsourcing some tasks to an MSP.

Deployment/Optimization

Deployment of VoIP is normally a staged process involving the following steps:

Deploy Network Enhancements

The modifications of the network identified in the network assessment are implemented and measurements are made to verify the expected improvements in network performance.

Simulate VoIP Performance

At this point it may be advisable to employ test applications or equipment that emulates voice traffic to verify network performance under simulated load conditions. This sort of simulated traffic testing can also verify QoS functionality and the effectiveness of the existing tools for monitoring and troubleshooting VoIP traffic.

Optimize for VoIP

QoS functionality is implemented throughout the network and IP header compression is configured to reduce overhead on voice payloads.

Pilot Deployment/Implement Full Deployment

In general, the transition to VoIP is a significant change to the network environment that justifies a pilot deployment followed by a phased production deployment, possibly on a department-by-department or site-by-site basis.

Management

As VoIP is undergoing its staged deployment, network management systems must be enhanced to extend the management model to include new VoIP devices deployed in the network, as well as the converged application traffic flowing over the network. The principal challenges involved in managing the converged infrastructure are:

Fault Management

Ensuring high availability of voice services requires minimizing mean time to repair (MTTR) for any faults that do occur in the underlying network infrastructure or in the VoIP devices. For example, 99.999% (or “Five 9s”) availability equates to only 5.25 minutes of downtime per year. Therefore, when failures do occur, the root cause must be rapidly identified irrespective of where the fault may occur in the end-to-end call path. An optimal fault management system can monitor all devices handling voice traffic, including the IP PBXs, VoIP servers/gateways, and TDM PBXs as well as the core IP infrastructure; i.e., switches, routers, firewalls, load balancers, etc. Managing voice calls across the VoIP and TDM boundary is particularly important considering the prevalence of phased VoIP deployments and the general inclination of the market toward hybrid VoIP/TDM PBXs.

Performance Management

Network evaluation and monitoring tools need to be capable of measuring the key end-to-end network performance parameters (delay, jitter, and packet loss) that are critical to the successful support of real-time traffic. Performance management tools that gather performance data for real-time and historical analysis are useful to help optimize the network design for real-time applications and for capacity planning purposes. In addition, it is highly advisable to deploy tools that automatically monitor communications sessions to detect in real-time when the levels of these key network parameters are trending to where they will soon exceed the targets for these parameters and then trigger an event or alarm. The management system should also be capable of monitoring call signaling traffic to measure the performance of the call control aspects of the VoIP system, such as delay-to-dial tone or frequency of call failure.

Management systems that provide integrated fault and performance management for both voice and data traffic are highly preferable because they maximize management effectiveness in dealing with service interruption, performance degradation as well as possible conflict between voice and data traffic. An integrated system also more readily supports automation of routine tasks such as traffic monitoring. Automation is important because it provides the earliest notification of impending problems or hard errors while freeing up human resources for higher value-added tasks.

Control

As previously noted, one aspect of controlling the simultaneous delivery of VoIP and data applications involves managing the different classes of VoIP and data traffic using rate-limiting QoS features and call admission levels for voice calls. Because of the dynamic nature of the enterprise application environment, frequent adjustments may be necessary to ensure that the goals for application performance continue to be met.

Another aspect of controlling the delivery of VoIP and data applications involves extending the security model to cover the converged network as well as the continual monitoring of the network for changes in traffic patterns that may have an impact on either VoIP or mission critical data applications. The basic goal in securing the converged network is to avoid the possibility of losing both data and voice communications due to a security event by preventing intrusions from spreading from the data environment to the voice environment and vice versa. This involves the logical isolation of VoIP and data traffic using separate virtual LANs (VLANs), plus deploying internal firewalls to safeguard IP PBXs and voice servers. Authenticating both VoIP endpoints and users to prevent intruders from using rogue devices to gain network access can also enhance the security of the voice environment.

CA Network & Voice Management Solution

A notable example of a management system optimized for fault, performance and voice management for converged networks is the CA Network & Voice Management (NVM) Solution. CA NVM Solution meets all the requirements for an integrated VoIP and data

management solution outlined in the previous sections of this document. CA NVM integrates the functionality of CA SPECTRUM® network fault manager, CA eHealth® network performance manager, and CA eHealth for Voice communication systems manager as if it's a single platform, accessible through a consistent user interface.

CA SPECTRUM

CA SPECTRUM provides network service, fault and configuration management across converged multi-vendor, multi-technology networks. Functionality includes extensive topology auto-discovery, relationship mapping, alarm notification, and configuration management. Patented event correlation and root cause analysis helps minimize MTTR by pinpointing degraded or failed network devices and suggests how the fault may be rectified. The impact analysis feature further determines what users and network services are impacted by various fault conditions or performance degradations.

CA eHealth

CA eHealth collects and analyzes real-time information to determine utilization patterns that violate usage thresholds or deviate from normal behavior. CA eHealth can feed into the CA SPECTRUM alarm system to leverage CA SPECTRUM's event correlation and root cause analysis functionality. CA eHealth can also provide real-time and historical data and reports required to document services levels, track bandwidth consumption, or plan for increased network capacity. CA eHealth functionality includes "what-if" analysis that can be used to model the impact of adding new requirements, such as VoIP, to the network. Therefore, eHealth can play a role in the planning of network upgrades to accommodate VoIP or to improve the level of service delivered for VoIP or critical data applications. Alternatively, this data can reveal what circuits, hardware or services can be downgraded or decommissioned to help cut operating costs.

CA eHealth for Voice

CA eHealth for Voice provides the ability to gather call data from both legacy and IP-capable PBX products from leading PBX vendors, including Avaya, Cisco, and Nortel. In addition to capturing historical data on voice traffic, CA eHealth for Voice and VoIP can continually test critical paths on the network to ensure that voice quality is not impaired by excessive delay, jitter, or packet loss. If established thresholds for these metrics are exceeded or abnormal behavior is detected, eHealth raises alerts to notify management personnel and can trigger automated fault analysis by CA SPECTRUM. Once probable cause has been identified by CA SPECTRUM, CA eHealth or CA eHealth for Voice can be used to obtain additional historical information to help confirm the diagnosis.

Summary

VoIP is the dominant approach that IT organizations are taking relative to new deployments of enterprise voice systems. VoIP offers a number of cost advantages as well as strategic advantages through the integration of voice communications with various data applications.

While the transition will take a few years to complete, it's possible to foresee a day when all voice traffic over enterprise networks (and even public networks) will be based on VoIP.

However, the addition of voice traffic to an enterprise IP data network poses a number of organizational and technical challenges. The key to meeting these challenges and successfully deploying VoIP is to follow a well-planned management process that has been proven successful by early adopters of VoIP and in the more general context of optimization of the performance of critical enterprise data applications over the LAN/WAN. As detailed in this white paper, some of the key steps in this process include:

- Model the impact of adding VoIP traffic to the network prior to deployment
- Implement the required network upgrades
- Monitor the ongoing performance of the VoIP traffic
- Track bandwidth consumption
- Map the VoIP traffic to the underlying network elements
- Identify degraded or failed network elements
- Quantify the impact of degraded or failed network elements
- Perform rapid root cause analysis

In order to be successful with the management component of this process, IT organizations must implement a network management solution that fully integrates fault and performance data from both the company's voice infrastructure and its IP network infrastructure. Integrated fault and performance management maximizes the quality of both voice and data services, improves service availability via reduced MTTR, enhances operator productivity, and facilitates automation of management processes. A notable example of such an integrated management solution for the converged network is the CA Network and Voice Management (NVM) Solution comprised of CA SPECTRUM, CA eHealth, and CA eHealth for Voice.

Bibliography

1. The 2006/2006 VoIP State of the Market Report, Steve Taylor, Distributed Networking Associates, <http://www.webtorials.com/abstracts/2005-2006VoIPSOTMReport.htm>
2. Enterprise telephony market tops \$9.6 billion in 2007, IP phone shipments up 29%, <http://www.infonetics.com/pr/2008/ms08.pbx.4q07.nr.asp>

3. The 2006/2006 VoIP State of the Market Report, Steve Taylor, Distributed Networking Associates, <http://www.webtorials.com/abstracts/2005-2006VoIPSOTMReport.htm>
4. The 2008 Application Delivery Handbook, Dr. Jim Metzler, <http://www.webtorials.com/abstracts/Kubernan2008handbook.htm>