

October 2004



Voice over IP Reliability

Architecture Matters



Comparing Three Different Approaches to Delivering 99.999% Availability

Phone service is like air in today's businesses: Its constant availability is both required and assumed. This state of affairs has arisen from more than a century of efforts by the makers of traditional voice switches to develop systems that are available at least 99.999% of the time. Concerns that voice over IP (VoIP) may not offer this vaunted "five nines" availability has been one of the major impediments to convergence.

This is ironic, because achieving a much higher level of communications-network resilience was the fundamental design goal for what is now the Internet. In contrast to their circuit-switched counterparts, packet-switched networks fundamentally enable fault tolerance, adaptive routing, and disaster recovery. It is quite possible—and can be quite cost-effective—to build IP-based voice systems that are more reliable than circuit-switched PBX platforms. The key is to start with the right foundation.

Today's VoIP solutions fall into three basic categories: Systems evolved from traditional PBX platforms, systems evolved from traditional data-switch platforms, and systems designed from the ground up for VoIP. All three of these architectures can be used to deliver VoIP systems with five-nines reliability, but they involve different degrees of complexity and cost. In this paper, we will examine the effects that these different architectures have upon the ability to deliver IP-based voice systems that are both highly available *and* cost effective.

The Legacy Voice Approach to VoIP Reliability

The manufacturers of legacy voice systems have a great history of delivering extremely reliable voice switches, and in fact are responsible for the 99.999% availability standard that VoIP solutions must match or exceed. However, legacy PBXes and key systems are hierarchical voice silos that operate independently at each location in a multi-site company. They cannot back each other up or be managed as a single voice network, and create a single point of failure at each site.

The VoIP solutions that have been evolved out of these legacy voice switches have inherited some of the inherent fragmentation of this centralized and hierarchical architecture.

One approach is to put a centralized IP PBX at the main site provisioning dial tone over an IP backbone to IP phones at the remote locations. If the WAN link goes down, the IP phones become useless, so availability assurance can be increased by installing a small standby IP PBX at each remote site. This device might be able to sense an outage and automatically take over as necessary, or it might require someone to flip a manual switch after employees start complaining that there is no dial tone, and then flip the switch back after the WAN link is restored. The cost of the failover solution is generally inversely proportional to its level of automation. If a company has multiple sites that are fairly large, separate IP PBXes are installed at each location.

The result is switches that operate more like separate silos than elements of a single voice system. This fragmented architecture cripples one of the key benefits of VoIP: the ability to create a single voice system that is distributed across multiple sites and can cover for individual switches that fail. It can also increase costs. Each of the silos may still require outsourced management and maintenance by a local teleconnect, just as their circuit-switched predecessors did. There may also be licensing issues to consider.

The Data-Centric Approach to VoIP Reliability

The term “convergence” sounds like a merger of roughly equal parties, but in the voice/data arena it is actually more of an acquisition. Voice is being turned into another network application, albeit a very important one. Given this fundamental truth, it is tempting to assume that the established data-switch vendors have an architectural edge as they develop VoIP platforms. However, the data world carries some legacy baggage of its own.

To begin with, while voice does not require much bandwidth, each conversation has to be maintained in a constant stream with strict latency thresholds for acceptable voice quality. Data traffic is very forgiving of latency, and the switches were designed to burst massive amounts of data through as quickly as possible.

The data-switch vendors approach VoIP by taking these basic, data-optimized Ethernet switch platforms and embedding some technology into them. Call control is implemented in a separate centralized server located in a data center. This centralized server is fundamental to the set up and tear down of all calls to, from, or between IP phones at all the company’s different sites, so it creates a single point of failure. Multiple call-control servers can be purchased at additional expense and clustered together to provide fault tolerance, but this centralized architecture still assumes that the IP backbone connecting the various offices is always available.

When the WAN does go down, the remote offices can fall back to a survival mode—if this optional feature was purchased and installed, and the IP phones were configured to use it in the event of a WAN outage. The approach is similar to that of the legacy voice vendors, who increase reliability by adding a standby PBX at the remote site. For the legacy data vendors, the enhanced availability takes the shape of a card that is installed in the remote site's data switch. This survival-mode feature may not be included in the quoted price of the VoIP solution, and it delivers a reduced level of voice service: Users limp along with dial tone and a few basic features.

When VoIP is based on a data-switch architecture, availability is increased by building in a lot of redundancies. This over-provisioning approach can get very expensive, and it also makes voice systems more complex. In fact, the inherent complexity of retrofitting a data platform for voice adds reliability challenges at every level: design, implementation, day-to-day operations, and problem resolution. Implementing a VoIP solution can involve sifting through hundreds of devices and options and trying to figure out which ones must be cobbled together to provide basic VoIP functionality with reasonable availability. Configuration of survival-mode features can require more than 20 command-line entries, each key peck adding one more opportunity for human error.

The VoIP-by-Design Advantage

IP networks are inherently distributed and resilient, and VoIP architects starting with a blank slate can exploit this fundamental strength to create a self-healing voice platform. A truly unified voice system can be distributed across multiple sites by using a simple peer-to-peer architecture that has no single point of failure. IP voice switches designed specifically for voice can each incorporate a complete call processor — even small models aimed at remote offices requiring eight ports or less. Each switch is a peer with a full complement of routing information safely held in local flash memory, and can operate as a standalone PBX if its site is cut off from the IP backbone. It can make best-effort calls on its own, using a failover PSTN trunk if necessary. When switches are added or restored to the network, they and the existing switches at all the sites automatically discover each other and start working together.

If a switch providing PSTN access to one site were to fail, its peer switches elsewhere in the WAN would provide alternate PSTN access to the users at that site. They would continue to get a full set of voice features, not a survival-mode subset. As long as the data backbone stays up, this type of distributed voice network can't have an outage unless all the switches go down simultaneously.

Reliability thus comes built in with this approach, and the five-nines availability requirement for voice is easily met. In fact, it can be increased to ten nines by installing a redundant switch with PSTN access at each site. A native VoIP architecture thus has built-in autonomy and survivability. Delivering a system that is reliable and highly available doesn't have to involve additional cost or complexity.

Conclusion

When disruptive technologies start to emerge, established vendors always look at the new landscape through the lens of the world they currently dominate. The architectures they come up with are limited or even crippled by this legacy. We all know that computers, generally speaking, have not been as reliable as phones, and there is still a lot of FUD (Fear, Uncertainty, and Doubt) circulating about VoIP. And in fact, there have been VoIP implementations that cannot be called reliable.

However, there is no longer any question that highly available VoIP systems can be delivered on any of the three basic architectures discussed in this paper. It is now a question of *how* such reliability is achieved, not whether it can be achieved. For VoIP solutions built on what is essentially a legacy voice or legacy data foundation, each additional nine in the availability rating comes at the expense of a lot more complexity and resources. For a voice system that is VoIP by design—mimicking the resilience of IP networks with a peer-to-peer communications architecture, and with autonomous call management built into each voice switch—it comes very naturally.



INTELLIGENT PHONE SYSTEMS™

960 Stewart Drive

Sunnyvale, CA 94085

(408) 331-3300

1-800-425-9385

Fax: (408) 331-3333

Email: info@shoretel.com

www.shoretel.com

© 2004 ShoreTel, Inc. All
rights reserved. October 2004