



Building Reliable IP Telephony Systems

**How Architecture and Design Differentiate
ShoreTel from the Competition**

**By Ed Basart, Chief Technology Officer
ShoreTel**

Table of Contents

1.0	Introduction	4
2.0	Five-Nines Availability	5
2.1	ShoreTel vs. Other IP PBXs	6
2.2	ShoreGear’s Demonstrated Availability	8
2.3	Predicted and Demonstrated MTBF	9
2.4	The Bathtub Curve	9
2.5	Mean Time to Repair (MTTR)	11
2.6	When More is More Trouble	12
2.7	Comparing Five-Nines Configurations	13
2.8	Five-Nines Availability Review	14
3.0	<i>N+1</i> Redundancy	15
3.1	No Single Point of Failure	16
3.2	Redundancy Review	16
4.0	Network Reliability	17
4.1	Network Reliability Review	19
5.0	Applications Reliability	21
5.1	Application Servers	21
5.2	Fallback Mechanisms	22
5.3	Server Effect on Availability	23
5.4	Wide Area Network Effect on Availability	25
5.5	Reliable Applications Review	25
6.0	Soft Reliability	26
6.1	Maintenance and Availability	26
6.2	Network Quality	27
6.3	Software Reliability Review	28
7.0	Conclusion	30
	Figure 1: ShoreTel ShoreGear Family	5
	Figure 2: ShoreTel Call Control Architecture	6
	Figure 3: Classic Chassis vs. ShoreTel Modular Design.....	7
	Figure 4: Serial Components Reliability Model	7
	Figure 5: Daughter Boards and Server PC Design	8
	Figure 6: The Bathtub Curve of Demonstrated Failure Rate.....	9
	Figure 7: Five-Nines: One ShoreGear vs. Competitor’s Multiple Units	14
	Figure 8: System with <i>n+1</i> Redundancy vs. 1:1 Redundancy.....	15
	Figure 9: <i>N+1</i> ShoreTel System Achieves Ten-Nines Availability	16
	Figure 10: Typical LAN Delivers Three-Nines of Availability.....	17
	Figure 11: ShoreTel Distributed Call Control Not Affected by WAN Failure.....	18
	Figure 12: ShoreTel Centralized Administration vs. Element Configuration	20

Table 1: Predicted Availability of a Typical ShoreGear Unit.....	5
Table 2: ShoreGear Availability	8
Table 3: Long Repair Times Slash Availability	11
Table 4: More Drives, Supplies and Fans Means More Service.....	12
Table 5: Multiple Servers Can Reduce Availability to Three-Nines	24
Table 6: Reliability Summary Comparison - ShoreTel vs. the Competition.....	29

1.0 Introduction

Reliability is the most critical aspect of a business phone system: You pick up the phone and you get a dial tone, period. ShoreTel delivers IP (Internet Protocol) telephony systems with unmatched reliability, using an approach that is fundamentally different from that of any other IP PBX supplier in the world. ShoreTel's architecture and design not only deliver high reliability, but do so in a very simple and cost-effective manner. This paper describes ShoreTel's unique approach to providing extremely reliable VoIP.

We start by defining reliability and availability, and then compare the approaches ShoreTel and other IP telephony vendors take to ensure high availability of their IP PBX hardware. You will see how the underlying system design and architecture dictate the type of redundancy that can be deployed to increase reliability, and why ShoreTel's $n+1$ redundancy is so much simpler and more cost-effective than the 1:1 redundancy used by other systems. The paper also addresses the reliability of the underlying data network and the challenges of implementing a virtually always-available voice system on an infrastructure that has a much lower availability rating.

No enterprise voice system today is without such applications as auto-attendant and voice mail, so we examine application reliability and the need to hold applications to the same reliability and availability standards as the voice system hardware. Finally, we finish by looking at soft reliability issues—the impact that software problems, administrative and maintenance activities, and network quality can have on voice system availability.

2.0 Five-Nines Availability

When voice system reliability is discussed, people are typically talking about the reliability of the hardware (as depicted in the figure of ShoreTel hardware, below). Without reliable hardware, you don't have a reliable system. We begin by defining hardware reliability and how it is achieved.



Figure 1: ShoreTel ShoreGear Family

Classically, reliability is measured by determining how often the hardware in a system fails and then computing the percentage of time the system is available. In telephony, the accepted benchmark is “five-nines” reliability, or a system that is available at least 99.999% of the time. We should note here that, while *availability* is what is actually computed, it is often mistakenly referred to as reliability, and spoken of as “five-nines of reliability.” Availability can be predicted based on the probability of hardware component failure, as detailed below.

Availability is predicted by taking into account the type and number of hardware components in a system and calculating the mean time between failure (MTBF). Currently shipping ShoreGear IP PBX units have a predicted MTBF of approximately 135,600 hours, and each failure requires one (1) hour of mean time to repair (MTTR). Using these variables, we can now do a simple computation to estimate the availability of ShoreTel hardware:

Table 1: Predicted Availability of a Typical ShoreGear Unit

$$\text{Availability} = \frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}} = \frac{135,600}{135,600 + 1} = 99.9993\%$$

This availability equation represents the standard definition of “reliability,” and indicates that a typical ShoreGear unit will achieve five-nines of availability. Stated another way, a ShoreGear switch on average is unavailable for one (1) hour every 10 years.

2.1 ShoreTel vs. Other IP PBXs

Other IP PBXs can achieve five-nines of availability, but they do so by using redundancy, which adds to the cost and complexity of the solution. The base units of competing IP telephony systems offer significantly lower availability than ShoreGear units because they lack the unique architecture and design of the ShoreTel system.

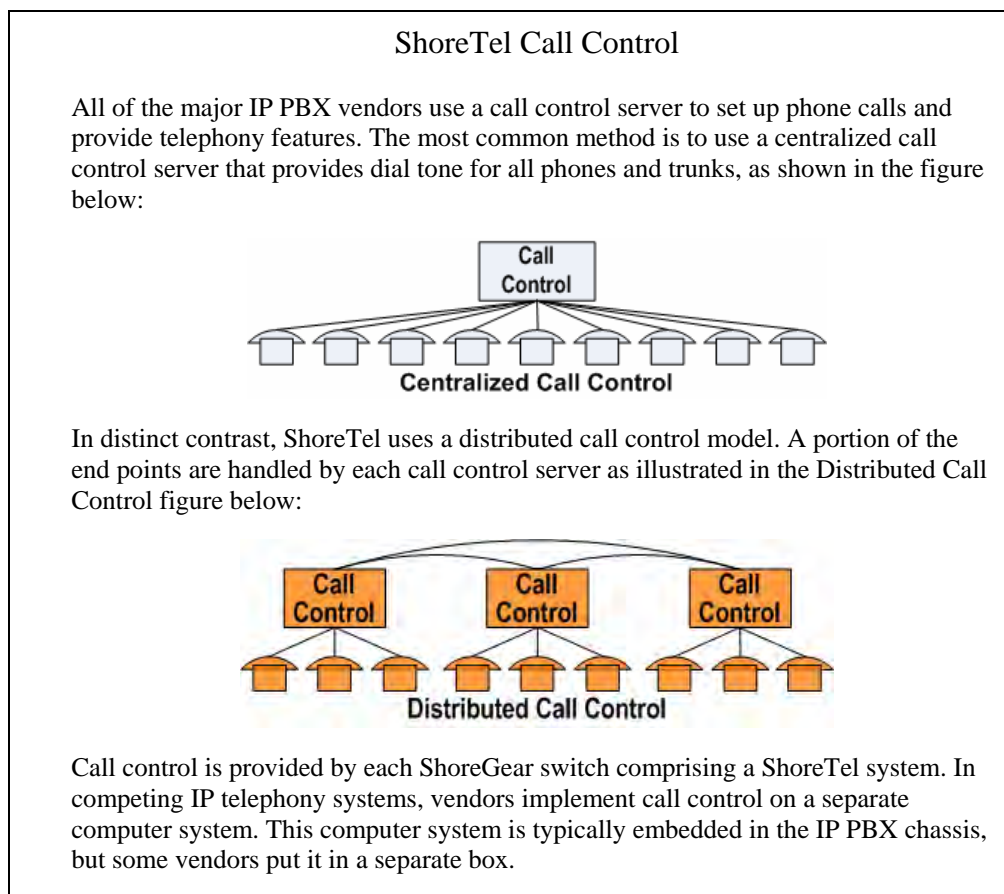


Figure 2: ShoreTel Call Control Architecture

The modular ShoreTel architecture allows for ShoreGear switches to be designed as small, simple, and reliable hardware. In contrast, most competitors use a chassis-based design of the type illustrated by the “Classic Chassis” figure below.




Figure 3: Classic Chassis vs. ShoreTel Modular Design

A classic chassis includes a number of circuit boards, with most providing telephony interfaces and one consisting of a specialized computer system. Compare this to ShoreTel’s modular ShoreGear units, which contain a single board. Because the ShoreGear modules have fewer and more reliable components, the ShoreTel system itself is more reliable. Again, reliability is calculated by adding up the failure rates of the individual components: The more components used, the lower the predicted reliability. This is due to the fact that all components must be working in order for the system to work.

Think of the typical telephony chassis as being like an old-fashioned string of serially connected Christmas lights. If one bulb fails, it takes out the entire segment. And the more lights on the string, the more vulnerable it becomes to such failures. This serial component model used in all major IP PBX systems is depicted in the figure below.



**Serial Components:
“String of Christmas Lights”**

Figure 4: Serial Components Reliability Model

Based on the reliability of their constituent components, such products typically have MTBFs in the 50,000 hour range, which translates to four-nines of availability. To get these chassis-based systems to five-nines of availability, the vendors use 1:1 redundancy, which is discussed below in the “*N+1* Redundancy” section.

Some vendors have come up with alternatives to the chassis-based system. One common approach puts telephony interfaces on daughter boards in a router and entrusts call control to a commercial PC server, as shown in the figure below. This architecture is similar to that of a chassis-based system, except that the addition of the separate server box further reduces reliability. How such an approach achieves five-nines of availability is discussed below in the “Comparing Five-Nines Configuration.”




Figure 5: Daughter Boards and Server PC Design

2.2 ShoreGear’s Demonstrated Availability

When designing a reliable product, you have to use expected MTBF numbers derived from an analysis of its constituent components, because the product itself doesn’t exist yet. However, once a product has been installed in the field in large numbers for a time, actual “demonstrated” reliability numbers may be used as an availability predictor.

Below are the predicted and demonstrated MTBF numbers for the ShoreGear family:

Table 2: ShoreGear Availability

ShoreGear Model	Predicted MTBF hours	Demonstrated MTBF hours	Demonstrated MTBF years	Availability, 1 hour MTTR
120/24	84,500	350,000	39.9	99.9997%
60/12	91,000	240,000	27.3	99.9996%
40/8	132,300	300,000	34.2	99.9997%
T1 and E1	154,200	350,000	39.9	99.9997%
Mean	135,600	330,000	37.6	99.9996%

All ShoreGear hardware units have a demonstrated availability that exceeds five-nines. This table also highlights three concepts we need to address at this point: demonstrated versus predicted MTBF; the ShoreGear 60/12’s anomalous demonstrated MTBF (the “bathtub curve”); and the Mean Time To Repair.

2.3 Predicted and Demonstrated MTBF

More than a decade of experience indicates that predicted MTBF is conservative, and that demonstrated MTBF numbers will be roughly twice as good. The method used by ShoreTel to calculate MTBF is the Bellcore TR322 standard, which was developed for traditional telephony equipment. Demonstrated MTBF hours are computed by keeping track of the number of units shipped over time, determining how many hours the units have been in service, and then dividing that total by the number of failed units returned for repair during that time interval. Because customer experience is based upon actual system performance, it tends to reflect demonstrated MTBF more than predicted MTBF.

The demonstrated MTBF numbers in the above table are for all ShoreGear units shipped between January 2003 and June 2006, and use return data through September 2006. As expected, ShoreTel's demonstrated MTBF is roughly twice its predicted MTBF, exceeding 300,000 hours.

2.4 The Bathtub Curve

Electronic products failures historically demonstrate a failure profile known as a “bathtub curve,” as illustrated in the chart below:

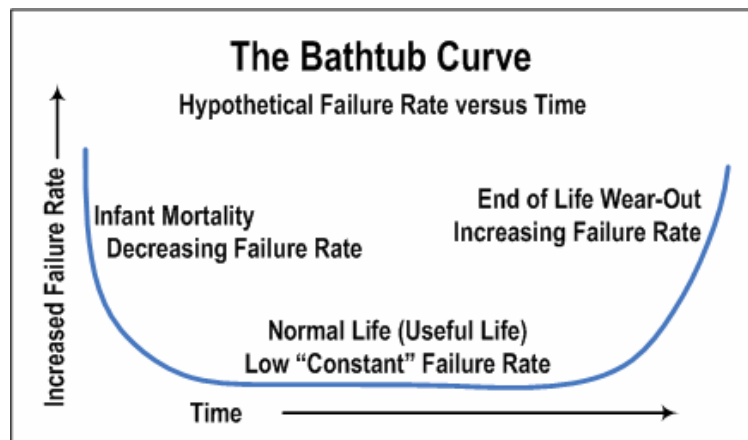


Figure 6: The Bathtub Curve of Demonstrated Failure Rate

For a number of complex reasons, electronics tend to fail early in life. Stress on the components is a major factor, and thermal stress plays a particularly significant role. After an electronic component is placed in service, thermal stress and other factors are responsible for some initial infant mortality before failure rates settle down at a much lower level. The ShoreGear 60/12 is the newest member of ShoreTel's ShoreGear family, and has only six quarters of field installations under its belt. Unlike the other ShoreGear family members, which have been in the field a year longer, the 60/12 has a

demonstrated MTBF that is less than double its predicted MTBF. We believe this is due to infant mortality still dominating failures, and expect that within 6 months, the 60/12's demonstrated MTBF will improve.

The other end of the bathtub curve illustrates end-of-life wear-out. The classic component subject to wear-out is the disc drive. Disc drive vendors recommend replacement usually after five years due to their complexity and the subtlety of their failures. For this reason, disc drives are not used in the ShoreGear family, as they cannot meet the requirements of a ten or twenty year life. In the ShoreGear family, the most significant contribution to failure is the fan (32%), although a long-life fan with a 500,000-hour MTBF is used. Fans fail primarily due to lubricant migrating away from the rotating bearings, ultimately causing the blade to stop moving because of the increased resistance.

Installing ShoreGear units in an air-conditioned room will reduce heat stress and minimize the impact of fan failures, thus increasing their demonstrated MTBF. The units contain both a fan rotation sensor and a temperature sensor, which generate alarms on the ShoreWare Director QuickLook administrative interface and enter events in the Windows event log. When a fan failure occurs, the system can notify the administrator by email, and the failed unit can then be swapped during the next scheduled maintenance period. Running ShoreGear hardware without a fan for long periods is not recommended, because the increased temperature stresses the components and accelerates wear-out failure. However, the tradeoff for tolerating short periods of increased wear is that down time can be scheduled and minimized.

ShoreTel reliability is not predicated on air-conditioning nor are we advocating replacing fans for preventative maintenance—we are just illustrating a component that is subject to wear-out on the bathtub curve and anecdotal mitigating factors. The power supply is another component that tends to follow the bathtub curve. Some capacitors “dry out” over time, and then begin to cause failures, particularly in the power supply. ShoreTel's power supply has a 550,000 hour MTBF. The fan and the power supply are by far the largest contribution to failure in a ShoreGear unit. Again, we do not advocate scheduling replacement of power supplies and fans, instead ShoreTel recommends “ $n+1$ ” redundancy—distributing the load across $n+1$ modules, where “ n ” modules are needed to carry the load, and should one module fail, the load is re-distributed among the remaining n modules. $N+1$ redundancy is the title and subject of the next major section, section 3.

2.5 Mean Time to Repair (MTTR)

The time that the system is down dramatically affects the system availability percentage. For example, substituting a 24-hour MTTR for the 1-hour MTTR in the availability formula above reduces system availability from five-nines to three-nines:

Table 3: Long Repair Times Slash Availability

One-Hour vs. 24-Hour ShoreTel Repair	
MTBF = 135,600 hours	
Availability = $MTBF / (MTBF + MTTR)$	
Availability with 1-hour MTTR	$= 135,600 / (135,600 + 1) = \mathbf{99.999\%}$
Availability with 24-hour MTTR	$= 135,600 / (135,600 + 24) = \mathbf{99.98\%}$

The more complicated the system, the longer it takes to identify the failing module, replace it, and restore service. Consequently, significant expertise is required to repair the complex chassis-based systems offered by other vendors, and a 4-hour MTTR is the standard for the industry. This in turn causes a problem for the IP PBX vendor that wishes to maintain five-nines of availability with a 4-hour MTTR. Typically, redundant systems are added on top of each other in an attempt to reach that esteemed five-nines rating. If the MTTR is 4 hours, a 400,000-hour MTBF is required to achieve 99.999% availability: $Availability = MTBF / (MTBF + MTTR) = 400,000 / (400,000 + 4) = 99.999\%$.

In contrast, ShoreTel's modular design makes system repair very simple, resulting in a much lower MTTR. Installation of the simple ShoreGear box requires only a power source and the connection of two or three cables. The administrative Director interface for swapping out a module is also very simple: Assign the MAC address of the failed unit to the new one, and all configuration parameters are automatically pushed into the box, so service can be restored quickly.

The simple ShoreTel modules also make it easy and cost-effective for a customer to maintain spares on site, as only two units are needed for a complete spare kit:

- 1 ShoreGear 120/24¹
- 1 ShoreGear T1 (or E1 internationally)

Products from other vendors require a large inventory of components. Keeping spare components or an entire spare chassis on hand is too expensive for most customers, and installing them requires more expertise than most customers command.

¹ Actually, customers need only match the largest ShoreGear IP PBX switch unit in their installation. The ShoreGear 120/24 can be dropped in to replace any model, although it may be overkill as a permanent replacement at smaller sites.

2.6 When More is More Trouble

Other IP PBXs use disc drives as a component of their server PC-based call control, while ShoreTel uses much more reliable flash memory. With their moving parts, disc drives have a much higher failure rate than flash memory, with demonstrated failure rates in the neighborhood of 500,000 to 1,000,000 hours. If you have an IP PBX unit with a 100,000-hour MTBF, adding a disc drive with a 500,000-hour MTBF to it cuts the system MTBF to 83,000 hours, or nearly in half.

Also, components with moving parts have important wear-out implications. The end-of-life portion of the bathtub curve for disc drives is very steep, and manufacturers recommend replacing the drive after its specified service life to avoid these high failure rates. Since most disc drives have a service life of 5 years, they have to be replaced during the lifetime of the IP PBX. Consequently, while the ShoreGear unit is an install-and-forget appliance, other vendors' products are more like a sailboat that requires constant attention.

Disc drives, power supplies, and fans are the most unreliable components of chassis-based IP PBXs, and vendors typically use redundant components to reduce the chance that any of them will fail. (But surprisingly a number of competitors have models without redundant components.) Ironically, while adding redundant disc drives, power supplies, and fans does allow the system to *survive* a failure, it also increases the chance that a failure will occur: The number of the system's least reliable components is being doubled. The following computation shows how these additions affect the system's mean time to required service:

Table 4: More Drives, Supplies and Fans Means More Service

Servicing Effects of Redundancy	
What happens when disc drives, power supplies, and fans—the system components with the lowest reliability—are doubled?	
Base unit MTBF	= 100,000 hours
Disc drive MTBF	= 500,000 hours
Fan MTBF	= 500,000 hours
Power supply MTBF	= 400,000 hours
Mean Time To Required Service (component failure)	= $1/(1/100,000 \text{ [base unit]} + 1/500,000 \text{ [disc]} + 1/500,000 \text{ [disc]} + 1/500,000 \text{ [fan]} + 1/500,000 \text{ [fan]} + 1/400,000 \text{ [power supply]})$ = 48,800 hours

Thus, the addition of two disc drives, two additional fans, and a second power supply slash the expected time before a repair is needed in half. The additional hardware may help with system survival, but it has the unwanted effect that repairs are now more frequent.

The data connectors used to incorporate cards into the chassis- and daughter-board-based IP PBXs provide an additional opportunity for trouble. Although predicted MTBF for a connector is no greater than many other components within the system, cards with connectors have a tendency to squirm because of thermal cycles or shipment-related shock. Eventually, corrosion occurs and intermittent problems start. The effects of intermittent connector failure is familiar to anyone who has shut down a multi-connector system like a PC to relocate it, discovered that it didn't work when re-powered, and then got it working again after taking it apart and re-seating the cables and boards. ShoreTel products have no data connectors and are immune to this problem.

To recap, the complexity of most IP PBX systems makes repair times longer, because diagnosing the problem, replacing the failed component, and restoring the system to service can be difficult. ShoreTel systems have a distinct MTTR advantage, thanks to simple modules which can be diagnosed and replaced quickly and easily.

2.7 Comparing Five-Nines Configurations

Now let us compare a ShoreTel IP PBX to that of a typical competitor. The figure below illustrates the configurations two systems require to achieve five-nines of availability. The ShoreGear unit on the left ships with five-nines availability, requiring no additional equipment. The competing system on the right implements the telephony components in a router (or possibly a gateway), and runs the call control software on a separate server PC. The server uses a disc drive, so a second drive is added for redundancy and to increase the availability of the PC. But even with a second drive, the PC falls short of five-nines availability, so a second PC with redundant disc drives is required. Also, note that the router and server PC must be connected via an Ethernet network, and since the Ethernet switch configured with redundant power supplies is not five-nines available, a second switch is required.

This configuration—recommended in the competitor's own reliability white paper—does achieve five-nines of availability. However, it takes a great deal more equipment, and three different types of equipment at that.

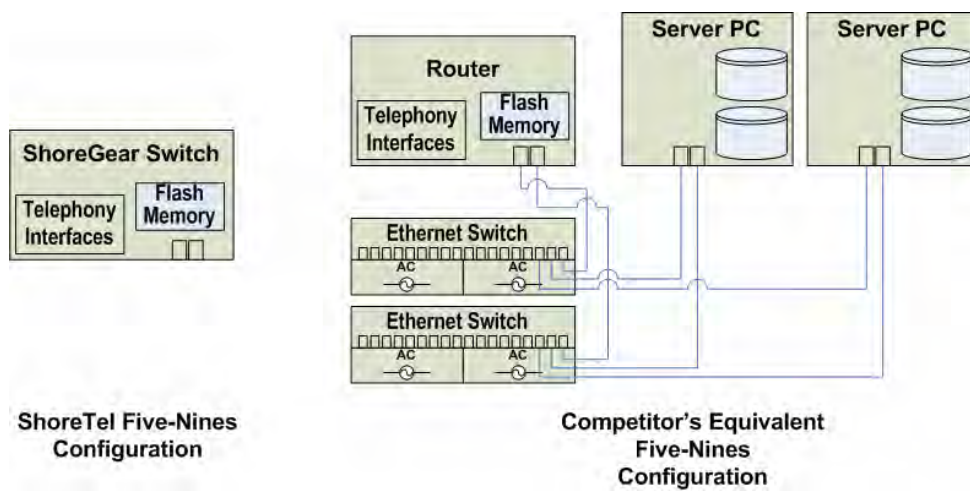


Figure 7: Five-Nines: One ShoreGear vs. Competitor's Multiple Units

2.8 Five-Nines Availability Review

1. ShoreTel hardware is very reliable, achieving five-nines of availability and more.
2. ShoreTel's architecture and modular design allow a low component count, making ShoreGear units inherently reliable.
3. ShoreTel uses no disc drives, avoiding a major reliability problem.
4. ShoreTel's modularity makes it easy to maintain spares and swap them in, reducing MTTR.
5. ShoreTel's competitors cannot achieve high reliability without redundancy, which increases cost and complexity.

3.0 N+1 Redundancy

ShoreTel’s uniquely distributed architecture allows the use of $n+1$ redundancy: Unlike competitors’ systems, which are all a “central box” design, ShoreTel’s model is to make each “box” a peer and put the intelligence in each box. The system load is distributed and shared across all boxes. This allows the system to scale linearly in size without forklift upgrades, as well as scale geographically across multiple sites. Using the ShoreGear units as an example, a typical system is implemented as n ShoreGear units. This is actually a form of redundancy (“ n ” redundancy). If an individual unit fails, “1/ n th” of the system is unavailable. The remaining users are unaffected.

Now if we add an additional ShoreGear to the system for redundancy, we have $n+1$ units, or $n+1$ redundancy. When one unit fails, the load is redistributed across the remaining units, and all users have use of the system. In contrast, ShoreTel’s competitors use 1:1 redundancy, in which a complete extra unit—twice the hardware—is needed. This comparison is illustrated in the figure below.

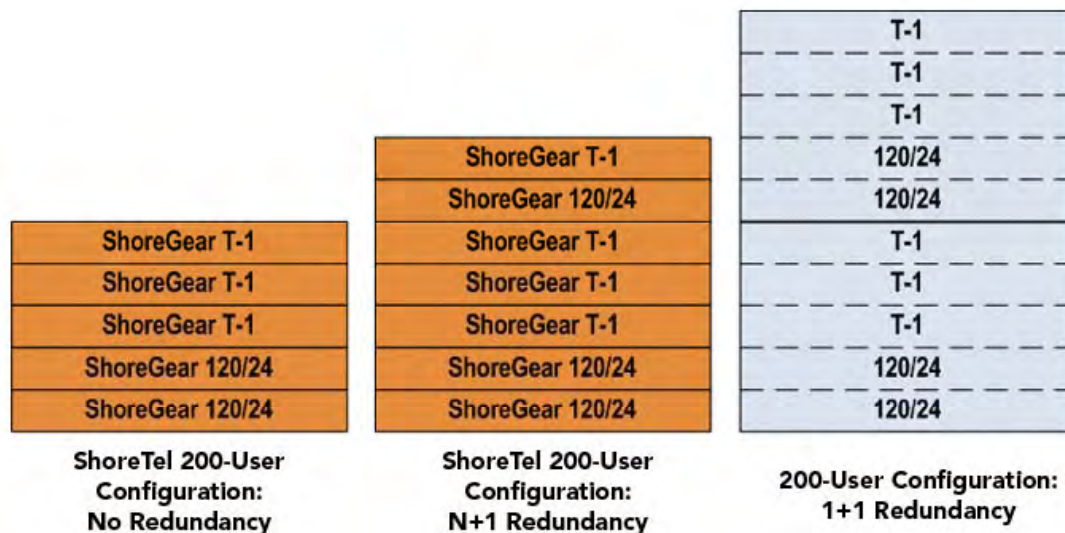


Figure 8: System with $n+1$ Redundancy vs. 1:1 Redundancy

Adding redundancy to a 200-user configuration requires two extra units for ShoreTel and five extra units for the competition. Also, all the units in ShoreTel’s $n+1$ redundancy configuration are functioning all the time. In the 1:1 redundancy model, the redundant unit is typically non-functioning, and standing by in “hot” (powered) or “cold” (un-powered) mode. In this inactive state, a failure of the standby unit can go undetected until it is called into service.

3.1 No Single Point of Failure

Using $n + 1$ redundancy creates a multiple-unit system with no single point of failure. If one T-1 in the 200-user redundant system above fails, the load is shared by the remaining T-1s. Similarly, if a 120/24 fails, the remaining 120/24s assume the tasks of the failed unit. The figure below illustrates the probability that the example system will fail. The T-1s and 120/24s each act as a parallel system, while the aggregate system is serial. As one can see, five-nines availability of each ShoreGear hardware unit results in an aggregate system with a miniscule probability of failure—ten-nines of availability.

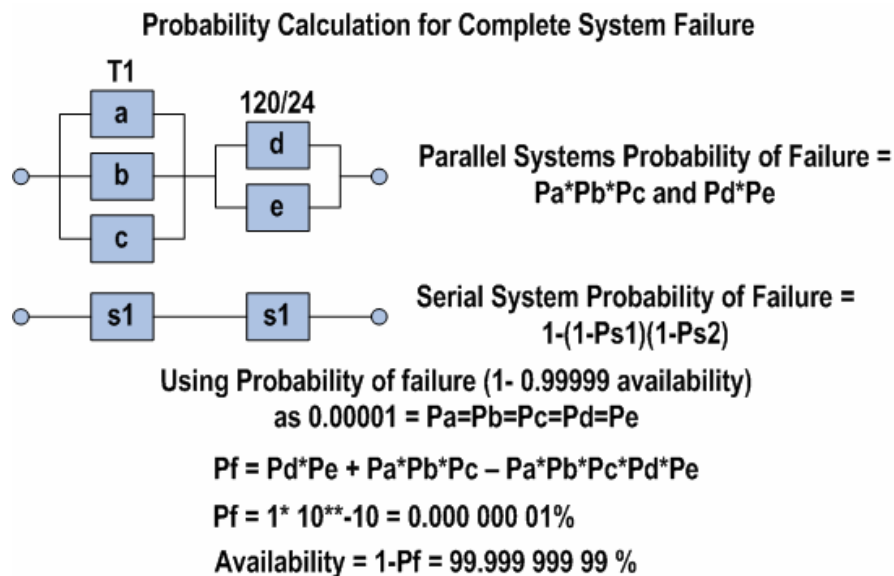


Figure 9: $N+1$ ShoreTel System Achieves Ten-Nines Availability

An $n+1$ system continues to be redundant, even after a single failure. As previously mentioned, the system is said to be n redundant. A failure of one of the n units degrades the system, and some services will be lost— $1/n$ th—while the remaining $n-1$ units are fully functional.

3.2 Redundancy Review

1. ShoreTel systems are inherently redundant, reducing the cost and complexity of reliability.
2. Competitors offer 1:1 redundancy rather than the more effective and efficient $n+1$ redundancy.
3. $N+1$ redundancy vastly exceeds five-nines availability.

4.0 Network Reliability

The most challenging aspect of an IP phone system is dealing with the underlying IP network infrastructure. Local area networks (LANs) and wide-area networks (WANs) have lower reliability than telecommunications systems and are prone to quality-of-service (QoS) issues that make IP telephony unreliable. Typical LANs achieve three- to four-nines of availability, as shown in the figure below:

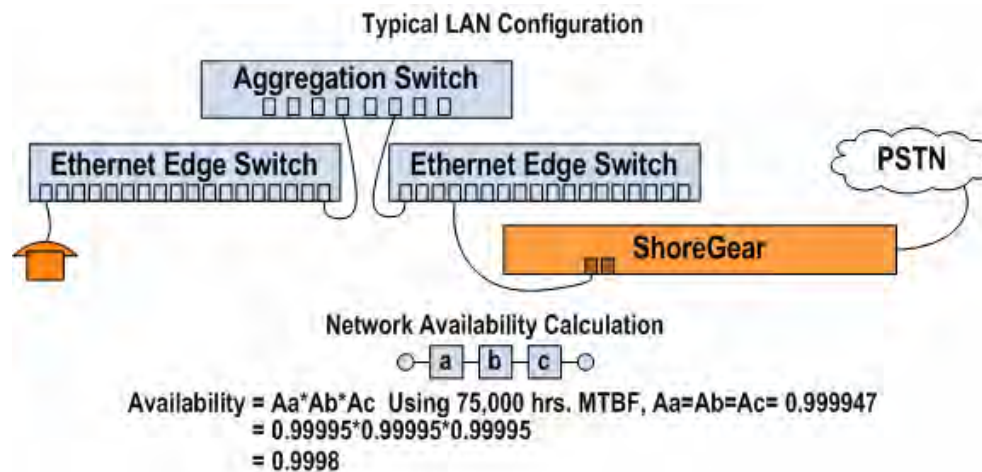


Figure 10: Typical LAN Delivers Three-Nines of Availability

LANs are less reliable primarily because they are implemented with multiple serial components. It is possible to achieve five-nines on the network by using a redundant aggregation switch that provides redundant paths to the ShoreGear IP PBX, although this does not appear to be a common implementation. It may be that data network implementers—who come from a technology culture much more forgiving of system failures and outages—are content with a four-nines system and an expected downtime of two (2) hours a year.

But reliable LANs are not the major challenge in networking, because they are a contained environment controlled completely by the individual enterprise. WANs are the real problem. WAN reliability numbers are not generally available, but our experience suggests that WAN links are available for basic connectivity 99% to 99.9% of the time, with voice-quality availability perhaps as low as 98%. The state of the art when the WAN is implemented entirely by a single service provider (which is best case) is 99.99%. The ShoreTel system mitigates this problem by distributing call control to each remote ShoreGear switch, along with a cached copy of the system database. When the LAN is down, each ShoreGear can continue handling calls, because call control, the business

logic, and the system database are all contained within the switch. ShoreTel’s distributed architecture has the “smarts” to provide seamless, full-featured operation despite WAN outages.

The figure below shows the difference between ShoreTel’s distributed architecture and the centralized systems of other vendors.

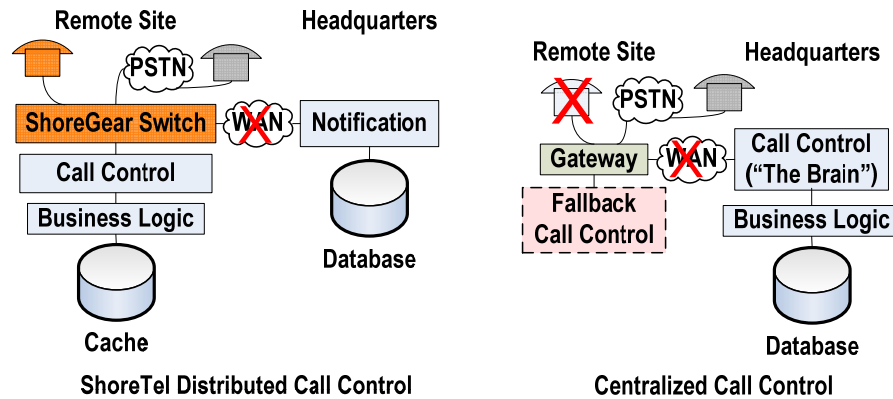


Figure 11: ShoreTel Distributed Call Control Not Affected by WAN Failure

The myriad features that make up a phone system are implemented by call control and business logic. Calls originate on end points and then are completed by call control, which consults business logic to come up with the correct destination and determine whether company policies permit the operation. The ShoreTel system maintains a central database at the headquarters site, and a notification service updates each switch as modifications are made. When the WAN goes down, each ShoreGear switch has the full set of features available.

A system with centralized call control is highly dependent upon its WAN connection. When the WAN goes down, the remote sites have no call control, and no calls can be made unless a fallback system has been installed. System implementers are stuck between a rock and a hard place: They can settle for fallback with reduced functionality, or implement full redundancy. The expense of the latter course makes this, in reality, a Hobson’s choice for most businesses.

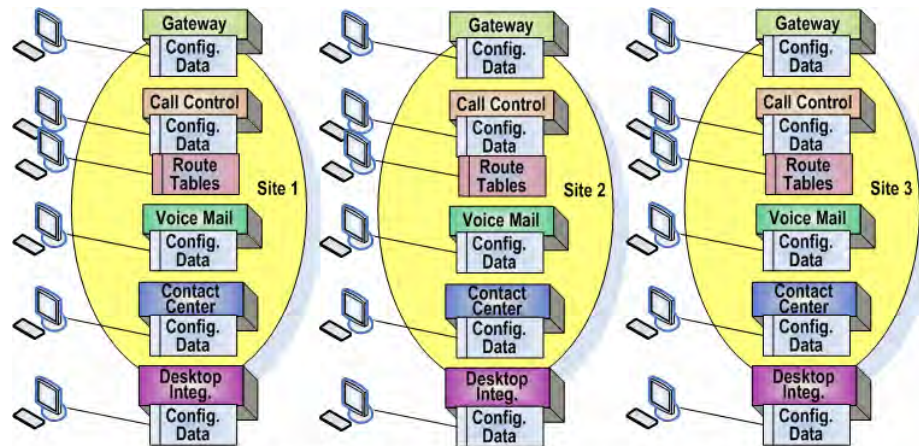
Simply put, other vendors’ products do not provide full, seamless call control functionality in the event of a WAN failure. To remain available when the WAN goes down, the voice switches at the remote sites are configured as independent PBXs. However, this requires tying the remote sites to headquarters with extra overhead. Also, it is difficult, if not impossible, to provide seamless call control functions across multiple PBXs. Refer to “Centralized Administration Versus Element Configuration” below for an illustration.

4.1 Network Reliability Review

1. *Data networks typically provide at most four nines of availability, making it difficult to support telephony's required five-nines of availability.*
2. *WANs typically provide only two- to three-nines of availability. Systems with centralized call control fail catastrophically without expensive add-on options.*
3. *ShoreTel's distributed architecture continues to provide seamless and complete call control at remote sites when the WAN fails, without extra cost or additional administrative overhead.*
4. *Other IP telephony systems provide full call control functionality at remote sites by using independent PBXs, thus losing the benefits of a single multi-site system.*

Centralized Administration vs. Element Configuration

Other products require element configuration. Each component has to be configured separately—and usually from a command-line interface, not a graphical one. As an example, consider the three-site system below. Each keyboard-and-screen icon indicates a separate configuration element: gateway, call control, voice mail and auto-attendant, workgroup call center, and desktop integration. And to tie the three individual IP PBXs together, route tables must be configured for each of the six elements at each of the three sites, which is another 18 configuration elements.



Site Element Configuration – Competitive Solution

In contrast, ShoreTel configurations are done from a web-based GUI that incorporates business logic to provide an integrated interface. All of the components shown above are administered with the browser-based ShoreTel, and ShoreTel’s Telephony Management Server distributes configuration data to individual components, as shown in the figure below:

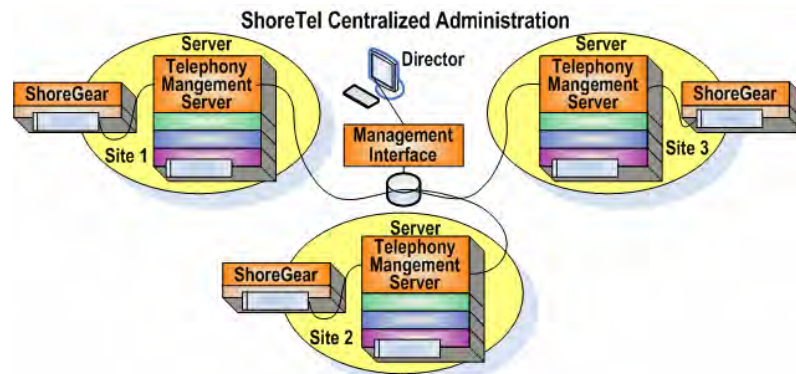


Figure 12: ShoreTel Centralized Administration vs. Element Configuration

5.0 Applications Reliability

The discussion so far has covered only traditional PBX functionality—the myriad call control features originally implemented in TDM-based voice platforms. When telephony reliability models were developed decades ago, applications such as auto-attendants, voice mail, and desktop integration did not yet exist, or were considered to be hopelessly exotic and expensive, and thus unessential in most environments.

When ShoreTel designed its first IP telephony system in 1996, (1) the reliability design objective was based on historical dial tone functionality requirements, and (2) these relatively new applications became integral, base-system components. When customers began actually using the ShoreTel system, expectations of high availability quickly increased to include the applications as well. We don't know if these expanded reliability requirements were just part of a rising tide of higher expectations by IP telephony users, or whether the ShoreTel system itself created them.

Traditional telephony systems were designed to deliver dial tone, and gradually started offering applications such as auto-attendant and voicemail as separate, add-on systems from third parties. Today many IP telephony vendors provide their own applications, but the same silo approach—multiple separate stacks of components and features—still prevails. This silo approach is illustrated in “Centralized Administration vs. Element Configuration” above.

5.1 Application Servers

ShoreTel systems may include more than one application server, but each application server includes a full set of applications:

- Auto-attendant
- Voice mail
- Application call control interface
- Call detail reporting (CDR)

In contrast, other IP telephony systems typically implement each voice application in a separate server, requiring multiple application servers—even in small installations.

A ShoreTel system is arranged as a site hierarchy, in which users are assigned to a site, utilizing the first application server above them in the hierarchy. The application servers each have full access to the configuration database maintained on the headquarters server, which is the root of the site

hierarchy. All application servers cache the configuration database in order to survive network outages or any other condition that makes the database unavailable to them.

Implemented in this way, the ShoreTel application servers are components in a modular architecture that enhances reliability. Each server provides applications for a fraction of the sites, and should it become unavailable, the outage affects only that portion of the overall system. The application server reliability features are outlined below:

- Auto-attendant. Each auto-attendant has a complete copy of all data, including recorded prompts, schedules, and menus. Users of the auto-attendant are unaffected by failure of other servers. Individual sites can be configured with multiple servers to provide auto-attendant backup should a server fail. When multiple auto-attendant servers are configured, they share the load, and a server is picked randomly to provide service.
- Voice Mail. When an organization deploys more than one voice mail server, each server contains the recorded name and personalized greetings for every user. If one of the voice mail servers fails, another will record messages for users assigned to the failed server, even if they are at a different site. When the failed server is restored, messages are transferred back to it from the fallback server. Individual sites can be configured with multiple servers to provide voice mail backup should a server fail. When multiple voice mail servers are configured, they share the load, and a server is picked randomly to provide service.
- Application Call Control Interface. ShoreTel users are provided with a desktop Call Manager application that integrates with Outlook and provides call control and visual voice mail. This application call control interface is basically unaffected by failure of other services; the Telephony Application Programming Interface (TAPI) endpoints controlled by the unavailable server are simply removed and then recreated when the server is restored.
- Call Detail Reporting. Each server collects CDR data at the end of each call and then forwards it to the headquarters server where all CDR data is merged into a single database. If the headquarters server is unavailable, up to two (2) hours of CDR data is cached locally and then transmitted to headquarters when service is restored.

5.2 Fallback Mechanisms

In an IP telephony system distributed over a network, there's always the possibility of a network outage that makes some portion of the system unreachable. When this occurs, the ShoreTel system

employs several fallback strategies that it employs to make the service degradation as graceful as possible:

- Failover Trunks. Each user is a member of a user group that can be configured with a prioritized list of trunk groups, with least-cost routing determining which trunk gets used in any given instance. Failover trunks can be configured for use in the event of network or hardware failure. All ShoreGear models have trunk interfaces that can be used for fallback, and they are often configured with a small number of analog trunks that take over when digital trunks become unavailable.
- Public Switched Telephone Network (PSTN) Failover. Should the IP network become unavailable when one extension calls another, the system will dial out on a trunk to the PSTN number for the extension—the Direct Inward Dial (DID) number.
- IP Phone Failover. Each ShoreTel IP phone maintains a heartbeat with its controlling switch. Should the switch become unavailable, the IP phone will automatically re-register with the system and be assigned another controlling switch.
- Switch Failover. Multi-site $n+1$ switch availability can be achieved with a single extra switch at the headquarters site that can cover for a failed switch at a remote site. Spare switches do not have to be maintained at each remote site.
- Backup Destination. When call control finds that an endpoint is unreachable, it attempts to use either the “forward no answer/busy” destination or the “forward-always” destination. This mechanism ensures that calls do not ring forever or “go nowhere” due to failures or mis-configuration.
- Backup Auto-attendant. If call control cannot reach an application server, each ShoreGear switch provides a backup auto-attendant which notifies the user that the destination is unavailable and offers to connect the user to another number.
- Copper Bypass. All ShoreGear switches have one analog trunk connected via mechanical relay to an analog extension in order to provide emergency service during a power failure.

5.3 Server Effect on Availability

ShoreTel implements its applications on one or more servers, and can distribute the applications among “ n ” servers, usually at multiple sites. We have chosen to call this form of redundancy n

redundancy (previously mentioned briefly in the No Single Point of Failure section). ShoreTel also bundles critical applications—auto-attendant, voice mail and desktop call control—onto each server, reducing system cost and complexity. The effects of n redundancy are discussed in the next section, Wide Area Network Effect on Redundancy.

Some ShoreTel competitors implement applications with a single server for each application. Using a server that is a single point of failure dramatically reduces overall system availability. If a functional phone system requires four major component services: 1) PSTN gateway, 2) call control, 3) auto-attendant and voice mail, and 4) desktop call control, then we can require that all four components must be available in order to say that the system is available. In some competitors' systems, each of these components are strung together in a serial fashion, with the resulting reliability of the previously mentioned old-fashioned string of Christmas lights—when one server fails, the system is down.

The availability of a server is approximately 40,000 hours, or four-nines. The overall availability can be roughly calculated, using five-nines for the PSTN gateway, and implementing each of the other components with a server. The availability is determined by multiplying together each of the availabilities, and as seen in the table below, the availability of the overall system is reduced from five-nines to three-nines!

Table 5: Multiple Servers Can Reduce Availability to Three-Nines

Adding a single server to a Five-Nines Base System:	
Availability	$\approx A_{\text{base system}} * A_{\text{server}}$
$A_{\text{base system}}$	$= 0.99999$
A_{server}	$= 0.9999$ (40,000 hours)
Availability	$= 0.99999 * 0.9999$
Availability = 0.9999	
Five-nines becomes four-nines	
Adding Multiple Servers to a Five-Nines Base System	
Availability	$\approx A_{\text{base system}} * A_{\text{call control server}} * A_{\text{voice mail server}} * A_{\text{desktop server}}$
Availability	$= 0.99999 * 0.9999 * 0.9999 * 0.9999$
Availability = 0.999	
Five-nines becomes three-nines	

Obviously, not all competitors rely on a string of single servers, since it is too unreliable. Competitors implement 1:1 redundancy on each server, and can thus maintain five-nines, but at increased cost and complexity, due to doubling the number of servers and having to manage more individual system elements and their interactions. As we will see in the next section, the 1:1 model becomes even more intractable in a multi-site network environment.

5.4 Wide Area Network Effect on Availability

Legacy Time-Division-Multiplex (TDM) PBXs were designed as single-site systems. A few best-of-breed implementations added an additional layer on top of the individual systems in order to approximate the effect of a single system. Such implementations were rare, costly and complex, and they were never able to achieve feature transparency from site to site. IP telephony has the promise to flatten a multi-site phone system into a single system—but only a fully distributed system such as ShoreTel’s can provide high availability in the face of three- or four-nine availability networks.

The competitors systems implement centralized call control at headquarters, and when the WAN is down, failover to a “survivable” call control at the remote site, with some reduced functionality. But none of the competitors offer applications redundancy at a remote site. This means applications have the same availability as that of the network—three- to four-nines. In addition, there is no real provision for redundant call control at remote sites for most competitors’ products. ShoreTel’s *n* redundancy for applications provides full functionality by installing a server at the remote site.

5.5 Reliable Applications Review

1. *Many applications are viewed as critical and also require five-nines of availability. ShoreTel bundles applications and holds them to the same reliability standard as the phone system itself. Some IP telephony vendors treat applications as separate silos from the system and in some cases provide no application redundancy.*
2. *ShoreTel has a number of fallback mechanisms for dealing with failures and making sure applications remain available, a critical requirement in a network environment.*
3. *Implementing a critical application on a server reduces the availability of the system to the availability of the server: four-nines. Implementing multiple critical applications as a series of servers reduces the availability of the system to three nines.*
4. *ShoreTel competitors do not distribute application servers to remote sites, as redundancy is limited to 1:1 at headquarters. This reduces application availability to that of the WAN—3- or 4-nines.*

6.0 Soft Reliability

Software issues are not commonly addressed in papers on reliability. There is no known metric for determining the impact of software reliability on system availability, and it is not a subject that has been discussed in vendor white papers. Nevertheless, we live in an increasingly software-driven world, and software issues loom large in the minds of enterprises shopping for IP PBXs. We thought about the metrics that ShoreTel's software organization maintains on bugs—such as bug severity, defects per line of code, mean time to fix a bug—but none of these measures capture how bugs affect the availability of the system to customers.

However, ShoreTel's support organization maintains an incident database that records all problems with the phone system as well as the nature and severity of the problem. A search of the database for problems that had a significant impact on service revealed 66 problems for approximately 3,000 customers during 2005. The records do not indicate precisely how long the problem persisted before it was resolved, but, based on 2006 data we know 4-hours is our average “open ticket” time for sever problems. We can then approximate the availability of the ShoreTel system including software failures:

$$\frac{\text{Available hours}}{\text{Total hours}} = \frac{100,196,880 \text{ hours} - 66 \text{ hours} \times 4}{100,196,880 \text{ hours}} = 99.9997\% \text{ availability}$$

We were surprised by this availability number, and we have started collecting and analyzing the information to better qualify both the concept of software availability as well as actual system availability to determine if ShoreTel is meeting the five-nines expectations. We have not yet collected enough data nor refined this system well enough to report our results, but hope to do so next year.

A simple cross-check can be made on the plausibility of the above availability. For hardware with 100,000 hour MTBF, there will be 8% annual failure rate. The 66 failures for 3000 customers is a 2% annual failure rate, which is roughly comparable to hardware failure and tends to validate the five-nines of overall availability.

6.1 Maintenance and Availability

Traditional hardware availability computations exclude maintenance periods. In reality, systems may be unavailable during maintenance while the following functions are performed:

- System shutdown time
- Software installation time
- Hardware removal/installation
- Reboots after installing line cards or gateways

In 24x7 environments, periodic administrative downtimes represent major problems, and ShoreTel keeps them to a minimum with its easy administration. Software distribution is handled by the system itself from the headquarters server, and the administrator can then selectively update on a site-by-site or switch-by-switch basis. Updates can be deferred until units are idle. This flexibility permits portions of the system to keep running while others are being updated, which avoids periods when the whole system is unavailable.

The ShoreGear switches themselves can be added or replaced without bringing the system down. And all system configuration changes (except one) can be performed without rebooting the switches. (The exception is the T-1 framer component, which requires a reboot if the signaling format is changed. Hardware reboots take approximately two minutes.)

ShoreTel's unique maintenance capabilities enable administrators to make major reconfigurations or upgrades and test them in as little as an hour. These same tasks often take 8 to 12 hours on other phone systems, which lack ShoreTel's simple browser-based GUI and must be administered by experts from an error-introducing command-line interface. ShoreTel owes this distinction to its simple design and architecture. Configuration changes are quick and easy because there is one single distributed system, regardless of the number of sites, and all changes—local or remote—are made from the same web-based ShoreTel Director console. In contrast, other IP telephony platforms are a series of disjoint systems that administrators must change individually by typing in text commands over and over again. This labor-intensive approach greatly increases the cost of providing a highly available system.

6.2 Network Quality

We should note here that IP PBXs, no matter how reliable, are very dependent upon the performance of the underlying IP network infrastructure. To meet the real-time requirements of voice, IP networks must be able to prioritize traffic and provide voice transmissions with a guaranteed Quality of Service (QoS). When the network fails to do this, a variety of utterly unacceptable voice errors—such as no dial tone, dropped calls, one-way audio, and “robot voice” call quality—start to occur.

Making their networks support such a real-time, mission-critical service is something new for system administrators from the more forgiving data world. However, configuring and maintaining a QoS network is outside the scope of this paper.

6.3 Software Reliability Review

- 1. ShoreTel's five-nines availability rating includes software failures*
- 2. ShoreTel can provide high availability even during system maintenance*
- 3. High IP PBX availability can be achieved only with a QoS network*

Table 6: Reliability Summary Comparison - ShoreTel vs. the Competition

Reliability Comparison: ShoreTel vs. the Competition				
Reliability Factor	ShoreTel IP PBX	Legacy PBX	Typical Enterprise IP PBX	Competitor Shortfalls
Architectural model	Distributed	Centralized	Centralized	One central point of failure
Scalability and backward compatibility	Scales up and down; still supports original units	Multiple products	Incompatible small system	Forklift upgrades; complex interconnects; multiple platforms to support
PBX style	Modular	Chassis	Gateway + server	Costly for smaller sites and hard limit for large systems
Physical construction	Single board	Classic chassis	Chassis or daughter boards	Increased component count lowers reliability, requires redundancy; mechanical components reduce reliability, may require expert repair
Call control memory	Flash memory	Disc drive	Disc drive	Lower reliability, requires redundancy; service life requires replacement
System reliability model	$N+1$	Redundant within chassis	1+1	Higher cost; single points of failure
Applications reliability model	Distributed	Redundant within chassis	Redundant within chassis	Higher cost; single points of failure
Fallback capabilities	PSTN, IP phone, destination, copper bypass, auto-attendant, voice mail, auto-attendant, desktop API, CDR	None	PSTN, IP phone	Failure is catastrophic
Software reliability	Only 0.35% of users per quarter make critical support calls	No public data	No public data	Lose repeatedly to ShoreTel in customer satisfaction surveys
Predicted maintenance time	<1 hour	Overnight	Overnight	Widely spaced software upgrades; big nightmares when they happen
System construction	Integrated; true single system	Silos	Silos	Disjoint product features
Management	Integrated; single-system view	Element by element	Element by element	Multiple management interfaces to master and use; overall system picture only in minds of individual managers

7.0 Conclusion

Nothing brings out the flaws in IP telephony systems quite like the need to deliver highly reliable and available voice service. Faced with architectural constraints, other IP telephony vendors offer clumsy solutions that leave enterprise customers grappling with difficult choices, lots of complexity, and expensive overhead.

In stark contrast, reliability and high availability are a natural outgrowth of ShoreTel's uniquely distributed architecture and elegant system design. The ShoreGear hardware is inherently reliable, delivering five-nines availability without embellishment. When multiple units are configured in a single distributed system, $n+1$ redundancy pushes ShoreTel availability above the traditional 99.999% benchmark by a significant margin, and at a lower cost than less effective solutions.

ShoreTel's distributed architecture continues to provide seamless and complete call control at remote sites that get cut off by WAN failures, without any extra equipment or administrative overhead. Availability ratings are not qualified in any way; edge components, applications, and voice-system software are all part of ShoreTel's five-nines-plus performance.

In short, ShoreTel takes a licking and just keeps on ticking, raising the reliability bar for both traditional and IP telephony 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

© ShoreTel, Inc. All rights reserved.