ShoreTel®
Brilliantly simple™

WHITE PAPER

Building Reliable IP Telephony Systems

How Architecture and Design Differentiate ShoreTel from the Competition

By Ed Basart, Chief Technology Officer, ShoreTel

Table of Contents

Abstract . . . 3
1. Introduction . . . 4
2. Five-Nines Availability . . . 4
2.1 ShoreTel Voice Switch's Demonstrated Availability . . . 5
2.2 Predicted and Demonstrated MTBF . . . 7
2.3 The Bathtub Curve . . . 8
2.4 Mean Time to Repair (MTTR) . . . 9
2.5 When More is More Trouble . . . 10
2.6 Reliability Comparisons . . . 12
2.7 Five-Nines Availability Review . . . 13
3. ShoreTel Distributed Architecture . . . 14
3.1 N+1 Redundancy . . . 15
3.2 No Single Point of Failure . . . 16
3.3 Distributed Architecture Review . . . 17
4. Network Reliability . . . 17
4.1 Reliable Remote Office Applications . . . 19
4.2 Network Reliability Review . . . 22
5. Fallback Resiliency . . . 22
6. Server Effect on Availability . . . 23
7. Soft Reliability . . . 24
7.1 Maintenance and Availability . . . 25
8. Summary and Conclusions . . . 25

Figure 1. The ShoreTel Voice Switch 220T1A . . . 4
Figure 2. The Bathtub Curve of Demonstrated Failure Rate . . . 8
Figure 3. Serial Components Reliability Model . . . 11
Figure 4. Centralized vs. ShoreTel Distributed Call Control . . . 14
Figure 5. Comparing PBX Hardware Designs . . . 15
Figure 6. System With N+1 Redundancy vs. 1:1 Redundancy . . . 16
Figure 7. N+1 ShoreTel System Achieves Greater Than Five-Nines Availability . . . 17
Figure 8. Typical LAN Delivers Three-Nines of Availability . . . 18
Figure 9. ShoreTel Distributed Call Control Not Affected by WAN Failure . . . 19
Figure 10. Typical Element Multisite Administration . . . 20
Figure 11. ShoreTel Distributed Applications . . . 21
Figure 12. ShoreTel Voice Switch "V" Models Provide Voicemail . . . 22
Figure 13. Servers Can Reduce Availability to Four-Nines . . . 24

Table 1. Predicted Availability of the ShoreTel Voice Switch 220T1A . . . 5
Table 2. Defining Reliability and Availability in Words . . . 5
Table 3. ShoreTel Voice Switch Availability . . . 6
Table 4. Long Repair Times Slash Availability . . . 10
Table 5. Redundant Components Require More Service . . . 11
Table 6. ShoreTel Reliability Comparisons . . . 13
Table 7. Reliability Factors, ShoreTel, PBX, IP PBX . . . 26

Abstract

The public telephone networks built decades ago have been incredibly reliable. Legacy enterprise telephony equipment can also achieve high reliability. IP telephony systems are now expected to delivery high reliability despite their newer technology, different implementations, and different features. This paper provides a tutorial that quantifies reliability, and describes industry techniques for achieving reliability, along with the difficulties encountered. The following are highlights of the topics covered:

- The meaning and measure of five-nines of mean time between failure (MTBF) reliability
- The contrast between delivering high uptime - availability - versus the statistics of reliability
- The superiority of n+1 redundancy versus traditional one to one redundancy
- No single point of failure as a guiding principal
- How networks provide the most difficult challenge to reliability
- Achieving five-nine availability with four-nine availability servers

ShoreTel's approach to reliability and its advantages are detailed in the following discussions:

- ShoreTel's 60-year demonstrated MTBF is approximately ten times higher than a typical PBX
- ShoreTel Voice Switch modular design is superior to chassis, daughter boards or server designs
- ShoreTel's use of n+1 redundancy vs. typical 1:1 redundancy
- ShoreTel distributed architecture solves lack of network reliability
- ShoreTel fallback adds resiliency

These aspects of ShoreTel's architecture and design deliver unprecedented reliability at low cost with brilliant simplicity.

1. 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 Internet Protocol (IP) telephony systems with unmatched reliability, using an approach that is fundamentally different from that of any other IP communication solutions 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 voice over IP (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. This 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 voicemail, so we also 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. Five-Nines Availability

When voice system reliability is discussed, people are typically talking about the reliability of the hardware (as depicted in Figure 1, The ShoreTel Voice Switch 220T1A, below). Without reliable hardware, you don't have a reliable system. We begin by defining hardware reliability and how it is achieved.

Figure 1. The ShoreTel Voice Switch 220T1A

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 percent 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." In this case, availability is a more accurate term compared to reliability, as the computation is 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, calculating the MTBF, include the down time for repair and convert to a percentage. For example, the ShoreTel Voice Switch 220T1A, depicted in Figure 1 above, has a predicted MTBF of approximately 163,500 hours. Assuming a one-hour mean time to repair (MTTR), we can now do a simple industry-standard computation to estimate the availability of this ShoreTel hardware, shown in Table 1, Predicted Availability of the ShoreTel Voice Switch 220T1A.

$$\text{Availability} = \frac{\text{MTBF}}{\text{MTBF + MTTR}} = \frac{163{,}500}{163{,}500 + 1} = 99.9994\%$$

Table 1. Predicted Availability of the ShoreTel Voice Switch 220T1A

This availability equation represents the standard definition of "reliability," and indicates that a typical ShoreTel Voice Switch unit will achieve five-nines of availability. Stated another way, the average time that a ShoreTel Voice Switch unit is unavailable is expected to be 32 minutes over ten years (0.0006 percent of ten years). See Table 2, Defining Reliability and Availability in Words, for definitions and distinction between the term reliability and availability.

Reliability:	The average time to hardware failure, showing how often it fails.
Availability:	The percentage of uptime - including repair - showing the combined effects of reliability and downtime.
Availability is the more inclusive measure.	

Table 2. Defining Reliability and Availability in Words

2.1 ShoreTel Voice Switch'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 does not yet exist. 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. See Table 3, ShoreTel Voice Switch Availability, for the predicted and demonstrated ShoreTel Voice Switch family MTBF values.

ShoreTel Voice Switch Model	Predicted MTBF hours	Demonstrated MTBF hours	Demonstrated MTBF years	Availability 1 hour MTBF
120/24 and 24A	84,500	530,000	60.5	99.9998%
60/12	91,000	528,000	60.2	99.9998%
40/8	132,300	520,000	59.3	99.9998%
T1 and E1	154,200	601,000	68.5	99.9998%
Half Rack Width Models				
90V	159,400	N/A	N/A	99.9994%
90BRIV	162,900	N/A	N/A	99.9994%
220T1A	163,500	N/A	N/A	99.9994%
90	171,400	215,000	24.5	99.9995%
30BRI and 90BRI	172,600	N/A	N/A	99.9994%
50V	175,800	N/A	N/A	99.9994%
220T1, 220E1, 220T1k, 220E1k	189,300	N/A	N/A	99.9995%
30 and 50	190,600	252,800	28.8	99.9996%

Table 3. ShoreTel Voice Switch Availability

All ShoreTel Voice Switches have a predicted availability that exceeds five-nines. All units with a demonstrated MTBF exceed their predicted MTBF, and also exceed five-nines. Not all ShoreTel Voice Switches have accumulated enough run time to calculate a meaningful demonstrated MTBF. Some of the half-width models — the ShoreTel Voice Switch 30, 50 and 90, began shipping in April, 2007, and their demonstrated MTBFs are shown. However, other models began shipping more recently, and have not yet accumulated a large number of run-time hours. Consequently these models' demonstrated MTBFs are shown as N/A (not available).

The demonstrated MTBF of 60 years is well above any other publicly published PBX numbers. One may come to the mistaken conclusion that a ShoreTel Voice Switch unit is guaranteed to last 60 years. Firstly, a 60-year MTBF means that approximately half of the units will have failed by the 60 year point (see Section 2.3 The Bathtub Curve for the failure rate over time). Units begin failing at intervals right from the moment they are installed. If we assume a constant failure rate, then nearly 2 percent of units will fail each year (if 60 years = 50 percent, 1 year = 1.7 percent). Secondly, a ShoreTel Voice Switch unit is expected to wear out at some point in time, with an expected 10-year to 20-year lifespan. Few units are likely to be working after 30 years of age. This principle is explained in more detail in Section 2.3 The Bathtub Curve.

This table also highlights three concepts we need to address at this point:

1) The meaning of predicted versus demonstrated MTBF

2) The "bathtub curve" — the ShoreTel Voice Switch 30, 50 and 90's anomalous demonstrated MTBF

3) The effect of the length of Mean Time To Repair (MTTR) on availability

ShoreTel
Brilliantly simple™

2.2 Predicted and Demonstrated MTBF

More than a decade of experience by the electronics industry shows that predicted MTBF is conservative, and that demonstrated MTBF may be twice as good, if not better. The method used by ShoreTel to calculate predicted MTBF is derived from the Bellcore TR322 standard, which was developed for traditional telephony equipment. The method used is called failures in time (FIT), which is the number of failures in one billion (10^9) hours of operation. This is usually interpreted as a large number of devices over a shorter time interval, e.g. 10,000 devices for 100,000 hours, or one million devices over 1,000 hours. It does not mean that a device is expected to last a billion hours; it's just the expected failure rate.

All of the individual FITs for each device component are added together, making the assumption that if an individual component fails, the device fails. The sum of all the FITs creates an expected FIT for the entire device. Then FIT is divided by one billion to produce MTBF (by definition MTBF is FIT divided by one billion). The model assumes that if any component fails, then the entire device fails. This assumption may not be true for many designs. For example, there are a large number of capacitors and resistors on the ShoreTel Voice Switch circuit boards. Many of these components exist for regulatory purposes and the unit can continue working after they fail, although it may have reduced margin or perhaps failed a regulatory requirement, such as electromagnetic emissions.

Two other assumptions about contributions to failure rate may not be true for ShoreTel, and that is the assumption that the devices may be power cycled a number of times and operate at high temperatures. ShoreTel Voice Switches are seldom powered down and generally operate in an air-conditioned room. These factors, plus the assumption that the device fails if any component fails, may be large contributors to the reasons why demonstrated MTBF is higher than predicted MTBF.

Demonstrated MTBF hours are computed by keeping track of the number of units shipped over time, summing the number of hours all 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 Table 3 are for all ShoreTel Voice Switch units shipped between January 2003 and December 2008, and use return data through March 2009. As expected, ShoreTel's demonstrated MTBF is significantly higher than predicted MTBF, exceeding 500,000 hours for a number of models. The ShoreTel Voice Switch 30, 50 and 90 models are newer models, and have lower demonstrated MTBF. This statistic is likely due to the "bathtub curve" described in the section below.

2.3 The Bathtub Curve

Electronic product failures historically demonstrate a failure profile known as a "bathtub curve" shown in Figure 2.

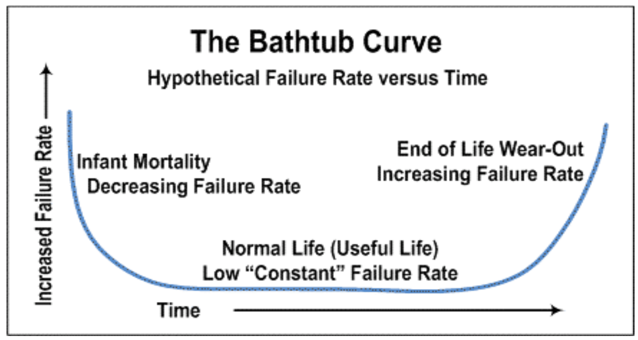

Figure 2. The Bathtub Curve of Demonstrated Failure Rate

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 typically settle down at a much lower level.

The ShoreTel Voice Switch half-width models have been in the field in volume for only about a year, and are demonstrating the infant mortality portion of the bathtub curve. We saw similar statistics as we introduced models in the past, and are confident that the half-width models will also provide superior demonstrated availability.

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 ShoreTel Voice Switch family, as they cannot meet ShoreTel's requirements for life expectancy.

ShoreTel Voice Switches use flash memory for storage. The expected life for a ShoreTel Voice Switch unit is ten years, and a unit might last as long as twenty years. We should note that flash memory now available for servers is much more reliable, and is beginning to replace disc drives. Currently, however, most servers still use disc drives due to the very limited size of solid state drives (SSD).

In the ShoreTel Voice Switch family, a significant contribution to failure is the power supply (17 percent of predicted failures for the ShoreTel Voice Switch 90) with an MTBF of 495,000 hours. Power supplies are dependent upon large capacitors, which "dry out" over time, and their failure rate rises significantly after ten years. They exhibit the end of life wear out portion of the bathtub curve. However, ShoreTel Voice Switch units are typically installed in air-conditioned rooms, keeping the units' temperature low, which in turn slows the drying out rate and extends the power supply life. This is another contributor to higher demonstrated MTBF, as discussed in Section 2.2, above.

Another ShoreTel Voice Switch component subject to the wear out portion of the bathtub curve is the fan, with an MTBF of 985,000 hours and the most significant contributor to failure (32 percent of predicted failures for the ShoreTel Voice Switch 90). Fans fail primarily due to lubricant migrating away from the rotating bearings and ultimately causing the blade to stop moving because of increased resistance.

Installing ShoreTel Voice Switch units in an air-conditioned room will reduce heat stress and minimize the impact of fan failures as units are likely to continue operating until downtime can be scheduled and the units replaced. The units contain both a fan rotation sensor and a temperature sensor, which generate alarms on the ShoreTel 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. Running ShoreTel Voice Switch hardware without a fan for long periods is not recommended because the increased temperature stresses the components and accelerates wear-out failure. However, the trade-off 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 scheduled replacement of power supplies or fans as preventative maintenance—we are just illustrating major components that are subject to wear out on the bathtub curve and anecdotal mitigating factors. 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 redistributed among the remaining n modules. N+1 redundancy is the subject of the next major section, Section 3, N+1 Redundancy.

2.4 Mean Time to Repair (MTTR)

MTTR can dramatically affect availability. For example, if you have installed a ShoreTel Voice Switch 220T1A with a spare unit on the shelf, in the event that the ShoreTel Voice Switch fails, the problem can be identified, the ShoreTel Voice Switch swapped and the system can be up and running within an hour. Achieving a one-hour MTTR along with an MTBF of 163,500 hours provides five-nines of availability. If you decide to have no spare, and have overnight delivery of a replacement, the expected time to repair is extended to 24 hours. This dramatically reduces availability from five-nines to three-nines as shown in Table 4, Long Repair Times Slash Availability.

One-Hour versus 24-Hour ShoreTel Repair
MTBF = 163,500 hours Availability = MTBF/(MTBF+MTTR)
Availability with 1-hour MTTR = 163,500/(163,500+1) = 99.999% Availability with 24-hour MTTR = 163,500/(163,500+24) = 99.985%

Table 4. Long Repair Times Slash Availability

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 four-hour MTTR is the standard for the industry.

This in turn causes a problem for the IP PBX vendor who wishes to maintain five-nines of availability with a four-hour MTTR, as a 400,000 hour MTBF is required for this to be achieved. (Availability = MTBF/(MTBF+MTTR) = 400,000/(400,000+4) = 99.999 percent.) Typically, redundant systems are required at additional cost and complexity.

In contrast, instead of plugging cards into a chassis, ShoreTel provides modular, independent ShoreTel Voice Switch units that can be racked and stacked. This design makes system repair very simple, resulting in a much lower MTTR. The administrative ShoreTel Director interface indentifies defective modules with a graphical "red light" in addition to providing failure notification with email and other types of events. Installation of the ShoreTel Voice Switch unit requires only a power source and the connection of two or three cables. The repair technician then uses ShoreTel Director to re-assign the configuration from the old switch to the new switch in one simple operation.

2.5 When More is More Trouble

All of ShoreTel's competitors' implementations have hardware that contains more components. Although conventional wisdom dictates that the larger the box, the more industrial and reliable it must be, a bigger box with simply more hardware reduces the reliability of the box. The reason more is not more reliable is that each additional component increases the chance of failure, as discussed previously with the FIT model.

For example, let's examine a typical competitive system with an MTBF of 71,000 hours (shown below in Table 5, Redundant Components Require More Service). The most unreliable individual components are the fan, power supply and disc drive. It is common to make each of these redundant in order to increase the availability of the system, but the increase is slight, improving only from 0.9996 percent to 0.9998 percent. What is not obvious is that three more of the most unreliable components have been added, and this actually increases the chances of a hardware failure, reducing the MTBF from 71,428 hours to 45,455 hours. To make the system five-nines available, the electronics must also be replaced, making a completely redundant system, but again reducing the MTBF to 35,000 hours.

	Typical System FIT	Redundant Fan, Power Supply, and Disc System FIT	Complete Redundant System
Electronics	6,000	6,000	2×6,000
Fan	2,000	2×2,000	2×2,000
Power Supply	4,000	2×4,000	2×4,000
Disc Drive	2,000	2×2,000	2×2,000
Total FIT	14,000	22,000	28,000
MTBF	71,428	45,455	35,714[1]
Availability (24 hr. MTTR)	99.96%	99.98%	99.96%[2]

MTBF to FIT	MTBF Hours	FIT (10^9/MTBF)
Electronics	166,667	6000
Fan	500,000	2000
Power Supply	250,000	4000
Disc Drive	500,000	2000

Table 5. Redundant Components Require More Service

The reason why adding components to a system reduces reliability is because MTBF calculations assume that all components must be working for the system to work. However, more components increase the chance of failure. Think of a (non-redundant) system 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 the string becomes to such failure. Figure 3 illustrates the similarity of electronic systems to a string of lights.

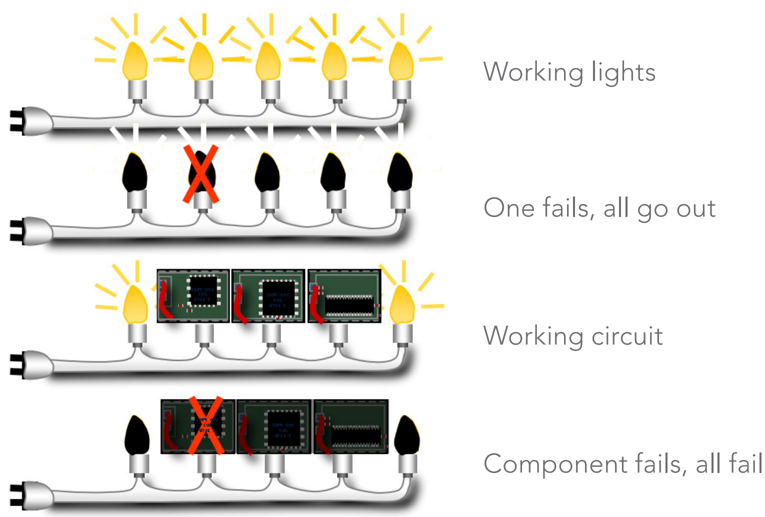

Figure 3. Serial Components Reliability Model

[1] MTBF went down due to redundancy additions. Availability went up as system is still working when a redundant component fails.
[2] Failure may be in a redundant component, with no system failure. However, the component must be repaired within 24 hours to meet the expected availability

All competitive PBX systems use complete redundancy for five-nines availability. Within the industry, a completely redundant system is dubbed "1:1" redundancy. An older telephony term, "tandem system," is also sometimes used. Contrast this with ShoreTel's system in which five-nines of availability is achieved in each individual ShoreTel Voice Switch, and redundancy is provided with n+1 redundancy (see Section 3, N+1 Redundancy).

Other factors also work against complex, chassis-style systems. The first problem is the tendency for higher heat (and higher power bills). Larger power supplies generate much higher heat, and are more prone to failure. Other components of the system also tend to be subject to higher heats and earlier failures.

Data connectors are used to incorporate cards and daughter boards into more complex systems and provide an additional opportunity for trouble. Although the predicted MTBF for a connector is no greater than other passive components, i.e. no chips or moving parts, connectors have a tendency to squirm because of thermal cycles and shipment-related shock. Eventually, corrosion occurs and intermittent problems start. The effects of intermittent connector failure are familiar to anyone who has shut down a multiconnector system like a PC to relocate it, discovered that it didn't work when repowered, and then got it working again after taking it apart and reseating the cables and boards. ShoreTel products have no data connectors and therefore are immune to this problem.

2.6 Reliability Comparisons

How does ShoreTel's reliability stack up against the competition? It's been difficult to collect published data on most devices. For critical systems such as PBXs and "carrier-grade" routers, vendors provide white papers describing how the product meets five-nines of availability, but do not publish calculated or demonstrated reliability numbers for those products. Frequently, there are MTBF numbers for individual components, such as power supplies and circuit cards, but insufficient information to calculate the overall system reliability in publicly available documents.

An Internet search reveals a number of white papers on reliability, and authors comment on the availability of routers, servers, and PBXs. These MTBFs vary from 30,000 to 70,000 hours. Disc drives are a frequent topic, with MTBFs from 500,000 hours to 10 million hours. Network reliability is discussed in papers by authors who have created highly reliable systems and estimated the availability of networks used by their systems. In Table 6, ShoreTel Reliability Comparisons, shown below, these and other devices and systems are grouped together by the number of nines of their availability[1]. Routers, servers and PBXs vary from four-year MTBF (40,000 hours) to less than a 40-year MTBF (400,000 hours). Disc drives, redundant PBXs and redundant routers come in at five-nines by necessity, as does an individual ShoreTel Voice Switch.

Network connections, such as DSL and T-1 WANs are significantly less reliable than any equipment at two-nines, and three-nines, respectively. Notice that although individual servers have a typical MTBF of 40,000 hours, when configured with 1:1 redundancy,

[1]Four hour MTTR used for all availability computations.

availability jumps dramatically from four-nines to well above five-nines. The ShoreTel n+1 solution availability jumps even more dramatically to far above five-nines. Both of these jumps can be arrived at intuitively by considering the probability of the system failing due to a second failure during the four-hour MTTR period after the first server or a ShoreTel Voice Switch fails. There is a very, very slim probability of this occurring. However, this does not mean in any way that redundant servers or ShoreTel Voice Switches will last for thousands or millions of years, as their expected wear out is in the 10+ year range. It is better to think of these large MTBFs as the probability of failure within a certain time period, which is shown in the last column.

MTBF Years	Item	Number of 9s of Availability	Probability of Failure in 1 Year
0.04	WAN Connection	2	1%
0.4	Router	3	0.1%
	Server		
	PBX		
4	500W PC power supply	4	0.01%
40	Disc drive	5	0.001%
	1:1 Redundant PBX		
	1:1 Redundant router		
	1:1 Redundant server		
	Individual ShoreTel Voice Switches		
4M	ShoreTel n+1 redundant ShoreTel Voice Switches	10	10^{-8} %
4B	The Sun becomes a red giant	13	10^{-11} %

Table 6. ShoreTel Reliability Comparisons

2.7 Five-Nines Availability Review

1. ShoreTel hardware has a predicted MTBF between 80,000 and 190,000 hours.
2. ShoreTel hardware has a demonstrated MTBF above 500,000 hours, or 60 years.
3. ShoreTel systems have an expected MTTR of an hour and availability well above five-nines.
4. Competitive systems tend to have lower MTBF, and achieve five-nines only through redundancy.
5. Complex chassis and rack-mounted systems have higher failure rates than simpler, non-redundant systems.

3. ShoreTel Distributed Architecture

Most IP PBXs are implemented with a central server running call control (sometimes called a call manager) to set up calls and provide telephony features. As we have seen previously, these systems use 1:1 redundancy to provide five-nines of availability. In contrast, ShoreTel uses a distributed call control model, in which each ShoreTel Voice Switch provides call control for a limited number of phones (and PSTN lines), and a system is implemented with multiple ShoreTel Voice Switch units. Figure 4 illustrates both central and distributed models.

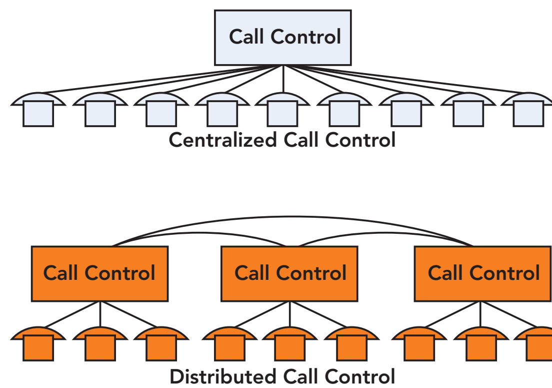

Figure 4. Centralized vs. ShoreTel Distributed Call Control

Competitive systems tend to implement call control on a standard server with disc drives, as compared to the ShoreTel Voice Switch, which is an embedded system. Servers are generalized computers, while embedded systems are specially designed for their task. Consequently, embedded systems tend to be more reliable and provide more reliable call control.

By distributing call control, ShoreTel Voice Switch units can be made smaller, simpler and more reliable as each unit takes on only a portion of the overall task. Other systems are larger; either "classic" chassis-based, or horizontal daughter board based, and sometimes use a separate server PC. All of these configurations are illustrated in Figure 5. The classic chassis contains a backplane, numerous connectors, large power supplies and fans in addition to the large number of components on each circuit board. The daughter board design eliminates the backplane, but sometimes requires adding cables, while retaining the connectors, and a number of components on each circuit board. The server PC has numerous cables, connectors, circuit boards and large power supplies.

As we have discussed previously, fans and large power supplies have a high probability of failure. The cables and connectors in these systems are frequently neglected in reliability computations, as it is difficult to assess mechanical reliability. But system administrators are familiar with having to re-seat cards and cables in these types of systems after moving them in order to get the equipment operating again. These are anecdotal observations, but reflect the actual problems of mechanical components in systems. Taking these reliability effects into account, along with the large component count in these configurations, their reliability is greatly reduced. Without redundancy, these systems are typically three-nines rather than the target of five-nines.

Because the ShoreTel Voice Switch units have a single board with fewer hardware components, no cables or connectors, few fans and small power supplies, the system is more reliable—a direct consequence of ShoreTel's distributed architecture.

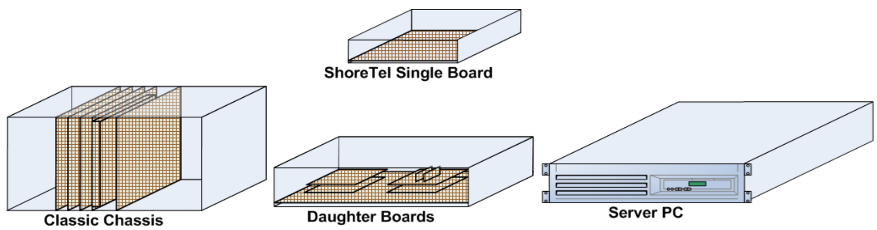

Figure 5. Comparing PBX Hardware Designs

However, despite ShoreTel's very high reliability, some form of redundancy is needed. Due to the increased number of components in larger systems with multiple ShoreTel Voice Switch, reliability may be reduced. ShoreTel's solution is n+1 redundancy, described in the following section.

3.1 N+1 Redundancy

ShoreTel's unique, distributed architecture allows the "unit" of redundancy to be much smaller. A ShoreTel system can be made redundant by the addition of an incremental unit, rather than duplicating the entire system. Under normal operation the load is spread across all units. When one unit fails, the load is redistributed across the remaining units. Adding one redundant unit to "n" units is known as n+1 redundancy, and ShoreTel is the only IP communication solutions provider to use this type of redundancy.

There are some additional benefits to the use of n+1 redundancy. Firstly, as there are multiple units, the system remains resilient even after a failure, as there are still n units remaining. The system degrades with each subsequent failure, but remains functioning for some users. Secondly, the units can be geographically distributed, making disaster recovery easier and redundancy much more cost effective for smaller, remote sites. A third benefit is that the system can scale without a forklift upgrade simply by adding ShoreTel Voice Switches.

Another important benefit is the incremental cost of redundancy. For an n+1 architecture, redundancy costs an extra 1/n. For an 1:1 architecture, redundancy costs 2n.

A comparison of a 500-user ShoreTel system implemented without redundancy is illustrated in Figure 6. Each ShoreTel Voice Switch 220T1 has a user capacity of 100, and in this example six ShoreTel Voice Switches are deployed in the basic system, and a seventh for n+1 redundancy. When a unit fails, users are automatically reassigned across the remaining six units. Should a second unit fail before the first unit is replaced, the system again reassigns users among the remaining five units. In this particular case, as the system was constructed with extra capacity, the five units have enough capacity for all 500 users, and there is no system degradation.

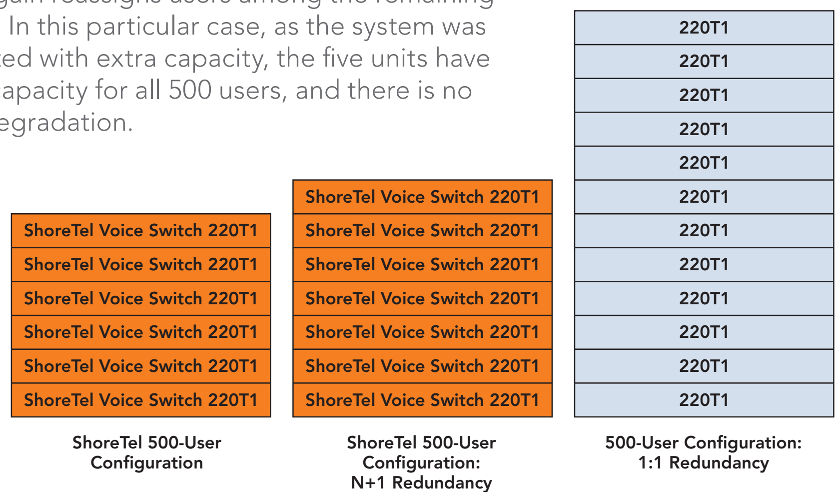

Figure 6. System With N+1 Redundancy vs. 1:1 Redundancy

3.2 No Single Point of Failure

An interesting approach for an infrastructure implementer is to follow the guidance of "no single point of failure." Naturally, highly reliable systems are desirable, but no matter how reliable the system, a failure for a mission critical system is intolerable. Of course no single point of failure means implementing redundancy, but the redundancy must be complete, otherwise a single failure can bring down the entire system. Here's how a single point of failure can be avoided. When a ShoreTel system is deployed, it is implemented as a number of modules - "n" in this case. A single additional module - "+1" - is deployed, creating an "n+1" system. The users are assigned across all the n+1 modules, and on individual module failure, the users on the failing module are reassigned among the surviving n modules. This means that an n+1 system has no single point of failure.

Taking the 500 user system above as an example we can configure as an n+1 system and calculate the availability as shown in Figure 7. Notice that the overall availability of the system with six ShoreTel Voice Switches and no redundancy is reduced from five-nines to three-nines (99.987 percent) This is due to treating all the units together as if they are in a series, and considering it a complete system failure if one unit fails. Of course this isn't strictly true as only one out of six users will be affected.

Adding the "+1" unit dramatically improves availability. The availability of an n+1 system is the probability of one unit failing and one of the n remaining units failing during the repair interval. In this example the availability of the overall system is significantly increased above five-nines, reaching seven-nines (99.999998 percent).

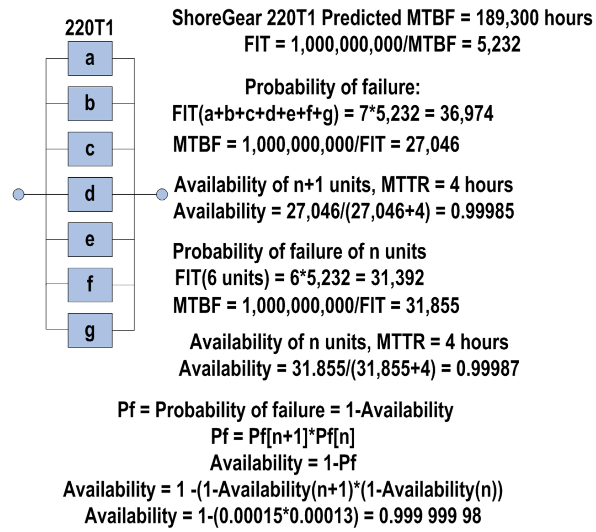

Figure 7. N+1 ShoreTel System Achieves Greater Than Five-Nines Availability

3.3 Distributed Architecture Review

1. ShoreTel's distributed architecture contributes to the reliability of ShoreTel Voice Switches by allowing units to be smaller with fewer components, and therefore more reliable.

2. ShoreTel's distributed architecture lends itself to n+1 redundancy, and is not confined to 1:1 redundancy.

3. N+1 systems are highly resilient, as a failure degrades the system rather than completely taking it out. Even without n+1 redundancy, a system implemented with n modules is highly resilient.

4. N+1 redundant systems can exceed five-nines availability.

4. Network Reliability

One of 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) are less reliable than telecommunications systems and are prone to quality-of-service (QoS) issues that can make IP telephony unreliable. In order for the ShoreTel system to maintain high availability, the LAN must be made redundant. Redundancy in WANs is more difficult and more costly than redundant WANs, and the ShoreTel system's distributed nature is used to overcome WAN issues.

Typical LANs achieve three- to four-nines of availability, as shown in Figure 8.

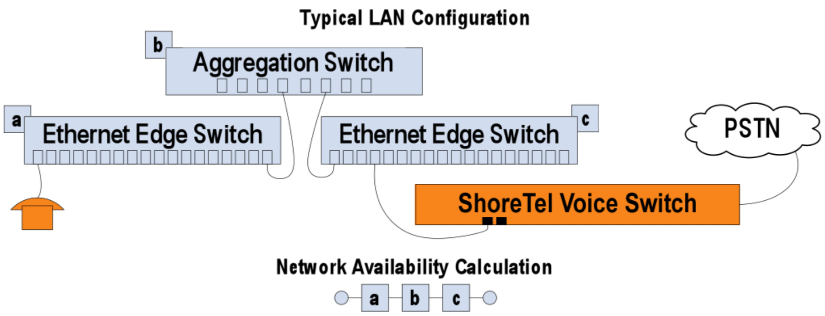

Using MTBF = 75,000 hours for a, b, and c. FIT = 10^9/MTBF = 13,333

Probability of failure: FIT(a+b+c) = 3*13,333 = 40,000

MTBF = 10^9/FIT = 25,000. Using MTTR = 4 hours

Availability = 25,000/25004 = 0.9998

Figure 8. Typical LAN Delivers Three-Nines of Availability

LANs are less reliable primarily because they are implemented with multiple serially connected components. It is possible to achieve five-nines on the network by using a redundant aggregation switch that provides redundant paths to the ShoreTel IP PBX, and is commonly deployed for large and midsized customers. Smaller customers tend to forego redundancy and rely in the inherent reliability of the equipment and a quick service call in case of an outage - content with a three-nines system and an expected downtime of eight hours a year.

Reliable WANs are the major challenge in networking, because unlike LANs, they are not under control of the individual enterprise, nor even the individual network provider. WAN reliability numbers are not generally available, but our experience suggests that WAN links are available for basic connectivity 99 percent to 99.9 percent of the time, with voice quality availability perhaps as low as 98 percent. The best of breed network providers promise 99.99 percent availability when the WAN is implemented entirely by that single service provider. There are no known network providers with a 99.999 percent service guarantee.

One possible solution to the WAN reliability problem is to purchase a redundant WAN. For large remote offices with IT staff, this is a reasonable solution. But for smaller offices, the cost of a second WAN connection that can carry voice traffic is cost prohibitive. Instead all major competitors either place a gateway at the remote site, or partition the system and install a small, separate PBX at the remote site. The competitive preferred solution is a gateway at the remote site, due to lower costs and less complexity. Figure 9 below shows both the ShoreTel distributed solution and competitive centralized solution for a remote site.

When the WAN fails the remote site loses contact with centralized call control — the single brain of the system. Phones at the remote site are placed in "survivable" mode and use a fallback call control located in the site's gateway. This means that the remote site loses the benefit of the detailed call control rules and permissions contained only at the central call control location. Call routing at the "survivable site" loses higher business logic, call permissions, and restrictions, and often resorts to functionality comparable to nothing more than the equivalent of a basic home phone line. In addition, any higher level telephony applications and services are unavailable. This loss of features and functionality and fallback to "nothing but PSTN access" is often not acceptable for many businesses that rely on telephony features and applications to run their businesses.

The ShoreTel situation is entirely different in the case of WAN failure. There is no change in call control as the local ShoreTel Voice Switch provides call control for the site, and continues to do so in the case of WAN failure. The call control in the ShoreTel Voice Switch maintains a cached copy of the database, and business logic. The database is maintained at headquarters, and change notifications are sent to all ShoreTel Voice Switch units in the system, while individual units update their database caches. The remote site acts as a fully functional, independent phone system when the WAN is down, and then transparently rejoins the main system when the WAN becomes available.

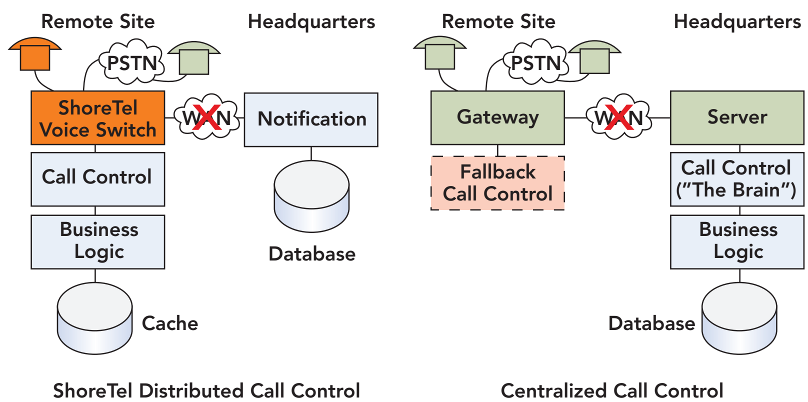

Figure 9. ShoreTel Distributed Call Control Unaffected by WAN Failure

4.1 Reliable Remote Office Applications

There are services needed at remote sites in addition to call control, and these are the variety of applications that run on servers. The most critical of these applications at the remote site are auto-attendant, voicemail, and in the case of ShoreTel, the desktop communication application ShoreTel Call Manager. In popular parlance, ShoreTel Call Manager is a unified communications application, providing desktop integration, presence, video and unified messaging.

The competitive solution to providing reliable applications for remote sites is to provide a complete system at these sites. The disadvantage is that users are presented with a disjointed feature set as the remote offices are implemented with a different product, even if provided by the same vendor. The independent phone systems are islands from a user perspective as well as from an administrative perspective. Competitive systems typically require separate administration of each of the major components of the system, imposing a heavy administrative burden as illustrated in Figure 10.

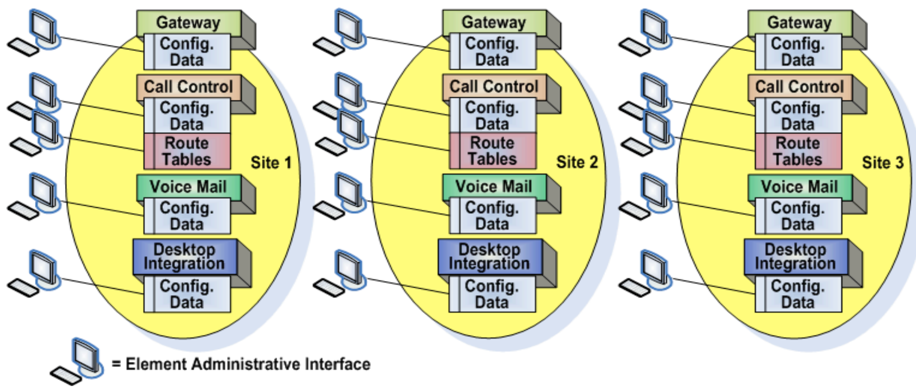

Figure 10. Typical Element Multisite Administration

Note that there are 18 elements to configure: six per system/site. Voicemail and desktop integration are replicated in order to achieve reliable applications. To stitch the separate phone systems together, dial plan and call route tables must be separately configured for system to system calling or long distance calling, adding cost and complexity to the solution.

On top of the complexity of a typical competitive multisite solution, all of the components at a remote site are a single point of failure, as 1:1 redundancy is too costly for small sites. This contrasts with the ShoreTel solution, which avoids these difficulties.

The ShoreTel approach is to embed application services in its distributed architecture. A Distributed Voice Services (DVS) server can be located at each site, as shown in Figure 11. The system maintains its single administrative interface – ShoreTel Director, and unlike competitive solutions, does not require duplication. Each server duplicates call control, voicemail, and desktop integration (the server side of the desktop application, ShoreTel Call Manager).

The Telephony Management Server coordinates the ShoreTel Voice Switches and application services on each server. The Telephony Management Server service notifies all the ShoreTel Voice Switches and services that are managed by this server of configuration changes and passes configuration data to them. Telephony application programming interfaces for server and desktop applications are also provided by the Telephony Management Server, allowing every application access to every phone. It also provides the "secret sauce" that ties all the phones, ShoreTel Voice Switches, servers and services together into a single phone system, also known as a single system image.

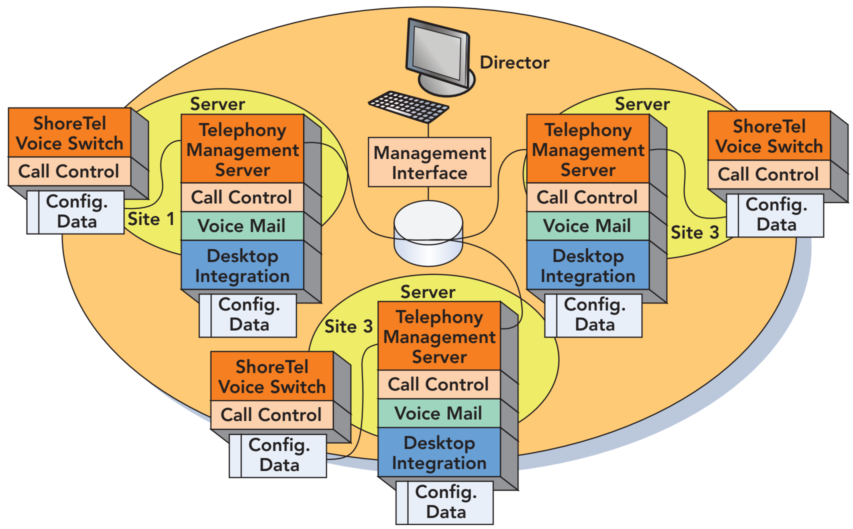

Figure 11. ShoreTel Distributed Applications

Two important redundant services are provided for remote sites: ShoreTel Voice Switch failover and voicemail and auto-attendant failover. A ShoreTel Voice Switch at the headquarters site can be designated as a failover ShoreTel Voice Switch for any or all sites. This is because protection is needed from a single point of failure. Should the WAN fail, that scenario is already covered by the ShoreTel Voice Switch and server on the site, and should the ShoreTel Voice Switch fail, it's unlikely that the WAN will also fail before the failed ShoreTel Voice Switch can be replaced.

In the case of server failure the ShoreTel Voice Switch unit finds an alternative voicemail/auto-attendant server. The servers are arranged in a hierarchy maintained by the system administrator. When the voicemail/auto-attendant server assigned to a ShoreTel Voice Switch is unavailable, the ShoreTel Voice Switch looks up the hierarchy until it finds a server and then uses it until its assigned voicemail/auto-attendant server becomes available.

For smaller sites, ShoreTel models with integrated voicemail eliminate the need for a server. These are the "V" series — ShoreTel Voice Switch 50V, 90V, and 90BRIV, illustrated in Figure 12, which provides voicemail/auto-attendant and flash memory storage. The same failover services are available. On ShoreTel Voice Switch failure, phones will be reassigned to a redundant ShoreTel Voice Switch and voicemail/auto-attendant will be provided by a server or series V ShoreTel Voice Switch found in the configured hierarchy.

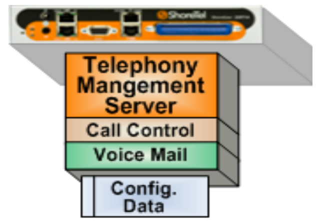

Figure 12. ShoreTel Voice Switch "V" Models Provide Voicemail

4.2 Network Reliability Review

1. Data networks have three- to four-nines of availability. To increase reliability at remote sites, organizations typically place equipment on site, rather than implement a redundant network.

2. Typical remote site implementations provide partial telephony functionality on failover. Placing a ShoreTel Voice Switch on site provides full functionality with no service interruption.

3. Typical remote site implementations provide no application services on WAN outages. The typical solution is an independent island phone system per site with reduced functionality, no redundancy and higher complexity. ShoreTel very simply maintains both functionality and redundancy at remote sites.

5. Fallback Resiliency

In an IP telephony system distributed over a network, there are a number of failure possibilities. Individual phones may come unplugged, the LAN may have a temporarily high error rate, the WAN may become congested, busy servers fail to respond, file systems reach capacity, and all equipment is subject to power failure. When any error or failure occurs, the ShoreTel system employs several resilient fallback strategies to make service degradation as graceful as possible:

- **IP Phone Failover.** Each ShoreTel IP phone maintains a heartbeat with its controlling ShoreTel Voice Switch. Should the unit become unavailable, the IP phone will automatically re-register with the system and be assigned another controlling ShoreTel Voice Switch.

- **ShoreTel Voice Switch Remote Failover.** Multisite n+1 ShoreTel Voice Switch availability can be achieved with a single extra unit at the headquarters site that can automatically cover for a failed ShoreTel Voice Switch at a remote site. Spare units do not have to be maintained at each remote site, which is a significant cost savings as the number of sites grows, reducing the cost of redundancy nearly in half.

- **Backup Destination.** When call control finds that an endpoint is unreachable, it attempts to use a configurable "backup destination." This ensures that calls do not ring forever or "go nowhere" due to failures or incorrect configuration.
- **Workgroup to Hunt Group Failover.** The ShoreTel Workgroup is commonly used as an informal contact center. In the event that the Workgroup service fails, calls can be sent to a hunt group to ensure that the call is answered.
- **Public Switched Telephone Network (PSTN) Failover.** If the WAN becomes unavailable, the system will dial out on a trunk to the destination PSTN Direct Inward Dial (DID) number to reach the remote site.
- **Failover Trunk Groups.** 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. A failover trunk group can be configured in the event of network or hardware failure. All ShoreTel Voice Switch models have trunk interfaces that can be used for fallback. Also used for fallback is a small number of analog trunks that take over when digital trunks become unavailable.
- **Backup Auto-attendant.** If call control cannot reach an application server, each ShoreTel Voice Switch switch provides a backup auto-attendant to notify the user that the destination is unavailable and offers to connect the user to another number to provide call answering.
- **Copper Bypass.** All ShoreTel Voice Switches with analog trunks have a copper bypass connected via mechanical relay to an analog extension in order to provide emergency service during a power failure.

6. Server Effect on Availability

Critical applications, such as voicemail and services for desktop applications, are frequently implemented on servers. In some implementations, it takes several servers to provide these applications. A typical server has four-nines of availability. As each of these servers is added to a five-nines redundant PBX, the availability quickly drops. Just two additional servers reduce the system availability to four-nines, as shown in Figure 13.

ShoreTel recommends implementing the system with multiple servers in over to avoid a single point of failure. Placing a server at each site and then failing over a service such as voicemail over to another site or to the headquarters server provides redundancy. The headquarters server can be backed up with an additional server, bringing the number of servers to n+1, modeling the n+1 availability model of ShoreTel Voice Switches.

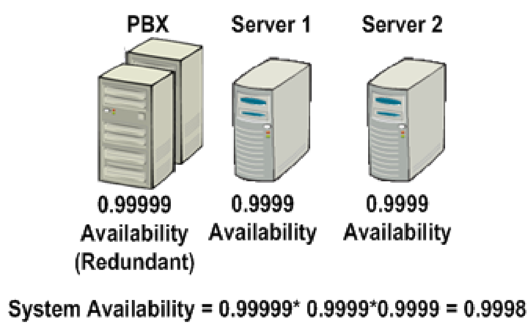

Figure 13. Servers Can Reduce Availability to Four-Nines

7. 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 much 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 for 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 227 problems with 279 hours of service-affecting impact for 6,741 customers during the last six months of 2008. We do not know whether the systems were actually down during this time period, but the customer considered the problem very severe, so we will treat the system as unavailable (which may be understating the system availability). The availability of the ShoreTel systems including software failures can be estimated:

$$\frac{\text{Available hours} = 28{,}889{,}472 \text{ hours} - 279 \text{ hours}}{\text{Total hours } 28{,}889{,}472 \text{ hours}} = 99.999\% \text{ availability}$$

We have been tracking this availability for the last four years, and have high confidence that five-nine reflect that overall ability of the system's software.

7.1 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 ShoreTel Voice Switches themselves can be added or replaced without bringing the system down. And all system configuration changes can be performed without rebooting the switches. The only exceptions are the T-1 or E-1 framer components, which require a reboot if the signaling format is changed. Hardware reboots take approximately two minutes to five minutes.

ShoreTel's maintenance capabilities enable administrators to make major reconfigurations or upgrades and test them in as little as an hour. These same tasks typically take much longer on other IP phone systems, which lack ShoreTel's simple ShoreTel Director single administrative interface and must be administered by experts from an error-prone 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 ShoreTel Director console. In contrast, other IP telephony platforms are a series of disjointed systems that administrators must change individually by typing in text commands over and over again. This labor-intensive approach greatly decreases the likelihood of maintaining a highly available system.

8. Summary and Conclusions

Table 7 summarizes the reliability factors presented in this paper, comparing ShoreTel with typical legacy PBXs and typical IP PBXs.

Reliability Factor	ShoreTel	Typical Legacy PBX	Typical IP PBX	Conclusion
Base availability	0.99999	0.9999	0.9999	ShoreTel Voice Switch is 5-9 vs. 4-9 availability
Demonstrated chance of failure (base system)	0.0002% (60 years/60 years + 1 hour)	Unpublished	Unpublished	ShoreTel Voice Switch is estimated to be nearly a factor of 10 times more reliable than a typical PBX
Mean Time To Repair (MTTR)	1 hour	4 hours	4 hours	ShoreTel architecture reduces MTTR and availability
Architecture	Distributed	Centralized	Centralized	ShoreTel architecture has no single point of failure
Call controll platform	Embedded	Embedded	Server	ShoreTel purpose-built hardware for highest reliability
Hardware design	Modular	Chassis	Chassis or daughter boards	ShoreTel modular construction increases reliability
Reduncdancy	N+1	1:1	1:1	ShoreTel redundancy has high availabilty and degrades rather than fails
Scalability	Single modular product	Multiple products	Multiple products	ShoreTel avoids forklift upgrade and insconsistent features across product lines
Design for WAN failure	Distributed to remote site	Separate system	Telephony survivable only, separate system	Full telephony critical applications
Remote site telephony redundancy	Failover to other site	None	None	ShoreTel sites have n+1 failover
Remote site voicemail	ShoreTel Voice Switch or server	Separate system	None or separate system	ShoreTel embeds in ShoreTel Voice Switch or server
Remote applications redundancy	Failover to other server	None	None	ShoreTel provides failover
Administration	Centralized	Element	Element	ShoreTel avoids error prone command line interface and individual elements
Resiliency	PSTN failover, phone failover, remote ShoreTel Voice Switch failover, backup destination, Workgroup failover to hunt group, failover trunk groups, backup autoattendant	Unknown	PSTN failover, phone failover	ShoreTel has a number of resilient strategies
Server redundancy	Similar to n+1 strategy	1:1	1:1	ShoreTel uses n+1 strategy, fewer servers
System software availability	Five-nines	Unpublished	Unpublished	ShoreTel tracks both hardware and software reliability and continuously improves
Maintenance downtime	One hour	Hours	Hours	ShoreTel central administration and distributed architecture shorten maintenance time

Table 7. Reliability Factors, ShoreTel, PBX, IP PBX

ABOUT SHORETEL

ShoreTel is the provider of brilliantly simple Unified Communication (UC) solutions based on its award-winning IP business phone system. We offer organizations of all sizes integrated, voice, video, data, and mobile communications on an open, distributed IP architecture that helps significantly reduce the complexity and costs typically associated with other solutions. The feature-rich ShoreTel UC system offers the lowest total cost of ownership (TCO) and the highest customer satisfaction in the industry, in part because it is easy to deploy, manage, scale and use. Increasingly, companies around the world are finding a competitive edge by replacing business-as-usual with new thinking, and choosing ShoreTel to handle their integrated business communication. ShoreTel is based in Sunnyvale, California, and has regional offices and partners worldwide. For more information, visit www.shoretel.com.

WORLD HEADQUARTERS 960 Stewart Drive, Sunnyvale, CA 94085 USA. www.shoretel.com
+1 (800) 425-9385 Toll Free. +1 (408) 331-3300 Tel. +1 (408) 331-3333 Fax.
EMEA +1 800 408 33133 Freephone. +44 (1628) 826300 Tel.
ASIA PACIFIC +61 (0)2 9959 8000 Tel.

ShoreTel®