# Ensuring Reliability in IP Telephony

Chapter 4

ShoreTel™

## TABLE OF CONTENTS

This chapter covers varying VoIP solution architectures, mean time between failure, mean time to repair, network reliability, and application reliability.

# Ensuring Reliability in IP Telephony

Chapter 4

The most crucial characteristic of a business phone system is reliability. You must pick up the phone to a dial tone, you must be able to successfully place outgoing calls, and calls must effectively reach your organization. This chapter covers varying VoIP solution architectures, mean time between failure, mean time to repair, network reliability, and application reliability. It is meant to help you dig deeper into the solutions you've narrowed down to your short list so that you can choose the one that fits best into your organization and existing infrastructure and provide you with maximum uptime.

## How is Reliability Different from Availability?

Usually when reliability is mentioned in terms of a voice system, the reference is generally about hardware. Without hardware reliability, the system cannot be reliable. Reliability is determined by calculating how often the system fails compared to the percentage of time the system is available. In the telephony world, "five-nines" reliability is the acceptable benchmark. This means the system is available at least 99.999 percent of the time.

Availability, on the other hand, is predicted based on the probability of a hardware component failure. It is predicted by taking into account the type and number of hardware components in a system and calculating the mean time between failure (MTBF). So, if an IP switch has a predicted MTBF of approximately 135,600 hours, and each failure requires one (1) hour of mean time to repair (MTTR), we would use this simple computation to estimate the availability:

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

This demonstrates that this particular unit will achieve "five-nines of availability." Alternatively, this switch is predicted to be unavailable for one hour every 10 years.

Let's take a household example. Consider a toaster that works for a year (an average year is 365.2425 days = 8,765.82 hours or 8,766 hours), and then it breaks, so you have to replace it: MTBF = one year. You take it to the store for a replacement the next day: MTTR = 24 (one day).

$$\text{Availability} = \frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}} = \frac{8.766 \text{ hours}}{8.766 \text{ hours} + 24 \text{ hours}} = 99.7\%$$

This indicates two-nines availability. However, if you keep an extra toaster on hand, MTTR could be as little as fifteen minutes (.25 hours). While this increases the cost of equipment, it also increases the availability fairly significantly.

$$\text{Availability} = \frac{8.766 \text{ hours}}{8.766 \text{ hours} + .25 \text{ hours}} = 99.997\% \text{ or four-nines of availability}$$

Back to industry terms, there is no ordinary telephone system that can achieve five-nines. Since state-of-the-art MTBF for systems is 100,000 hours and MTTR is 24 hours, you would need to deliver 2,400,000 hours between failures to achieve five-nines.

$$\text{Availability} = \frac{\text{MTBF}}{\text{MTBF} + 24} = \frac{2,400.000}{2,400,000 + 24}$$

Even repairing the problem in 4 hours doesn't make it much easier to accomplish:

$$99.9990\% = \frac{400,000}{400,000 + 4}$$

You would still need 400,000 hours between failures. These examples are far beyond state-of-the-art. The way to meet these demands is via redundancy. Read on for a section on redundancy and specifically n+1 redundancy.

## Distributed vs. Centralized, Chassis vs. Modular

IP telephony systems differ in their architectures: Some are centralized while others are distributed. In a centralized setup, the centralized call control server provides dial tone for all phones, whereas a distributed model is one where end points are handled by multiple call control servers. In this solution, call control is provided by each switch in the system. See figure 4.1.

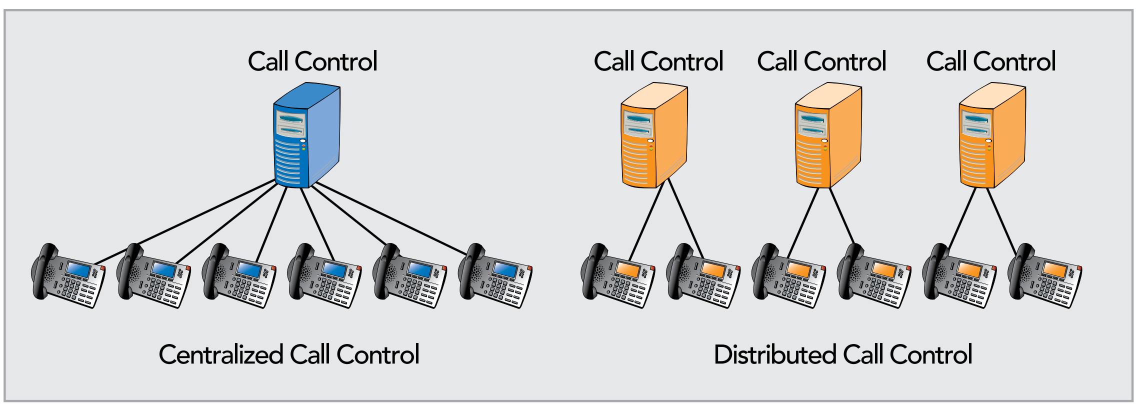


Figure 4.1. Centralized vs. Distributed Call Control

A classic chassis includes a number of circuit boards, with most of them providing telephony interfaces and one consisting of a specialized computer system, while some modular units contain a single board. The classic chassis can be compared to a string of holiday lights: If one bulb fails, the entire segment fails. The more lights on the string (number of circuit boards in the chassis), the more vulnerable it becomes to failures.

A typical chassis model, because you have to take into consideration the reliability of their components, typically has an MTBF in the 50,000 range, which is four (not five) nines availability. This can be raised to five-nines by adding switches for redundancy (costly but effective). More on this will be discussed in the n+1 redundancy section later.

In contrast, a modular architecture includes small, simple and reliable hardware. This modularity is more reliable and also offers more freedom in the design stages of a VoIP implementation. Look at both modular and chassis-based systems, but keep in mind your specific reliability needs and remember that modular systems generally make configuration changes simpler and seamless.

## The Bathtub Curve

Electronic product failures historically demonstrate a failure profile known as a "bathtub curve." See Figure 4.2 for a depiction of the bathtub curve. Because of a number of reasons, including stress, electronics tend to have a short life before they start failing. At the beginning of the lifecycle (the left side of the diagram), manufacturing defects, defective parts, contamination and other factors cause failures, before these settle to a much lower level (the middle of the diagram). The other end (on the right) signifies the end of life or wearing out of the product.

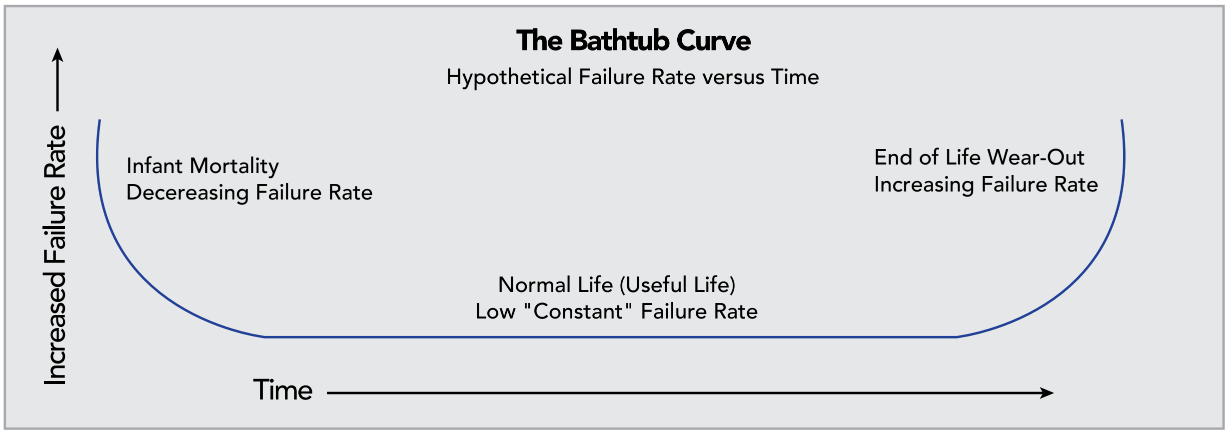


Figure 4.2 The Bathtub Curve

Be sure and ask vendors about their failure rates and how long a product lasts before end of life. If a vendor does not give you a concrete number based on scientific calculations (not marketing hype), ask more questions or talk to someone at the organization who can give you that information.

## Mean Time To Repair (MTTR)

When a product is down, the entire system's availability percentage is dramatically affected. Consider the following example, where MTTR goes from 1 to 24 hours.

**Long Repair Times Slash Availability Ratings**

MTBF = 135,600 hours
Availability = MTBF/(MTBF+MTTR)

Availability with 1-hour MTTR = 135,600/(135,600+1) = 99.9993%

Availability with 24-hour MTTR = 135,600/(135,600+24) = 99.9789%

Figure 4.3 Comparing MTBF with varying MTTR

The more complex an IP telephone system, the longer it's going to take to identify what's going wrong during a failure. Only when you've identified what's wrong can you get a replacement for it, which can take even more time, and then there is the time it takes to get the system back up and running. Because of this, chassis systems described earlier in this chapter require personnel with significantly more expertise to ensure the system remains functional.

A 4-hour MTTR is industry standard, which creates a problem for IP vendors that want to maintain five-nines of availability with a 4-hour MTTR. Redundant systems are usually added to ensure this availability because a 4-hour MTTR requires a 400,000-hour MTBF to achieve 99.999% availability. (Availability = MTBF/(MTBF+MTTR) = 400,000/(400,000+4) = 99.999%.) Modular, distributed systems tend to make system repair easy, which results in a lower MTTR. These systems only require one power source and two or three cable connections.

## Moving Parts and Complexity

Another thing to keep in mind is the number of moving parts there are in a system. For instance, adding a disc drive (rather than flash memory) with a 500,000-hour MTBF cuts the system's overall MTBF in half. Moving parts are also likelier to wear out faster than non-moving parts. For instance, the bathtub curve for disc drives is steep and it's often recommended they be replaced well before end of life to avoid failure. In the case of an IP telephony solution, you'd be replacing a disc drive during the time you have it on your network, since most disc drives last five years. Ask each vendor how many moving parts there are in each system. Again, insist on getting the information from another company source if the sales team does not have this information readily available.

Redundancy also impacts the failure rate, ironically. While vendors often add redundant parts, such as disc drives and power supplies, to their systems, the very fact that the number of parts are being doubled in itself can increase the chance that the system will fail (increase the MTBF). When you are considering an IP PBX system for your organization, be sure to look at how complex each system is. The more complex, the longer it takes to repair because problem diagnosis, part replacement, and system restoration can be difficult. Look for modular systems that are easy to manage and troubleshoot, with specific built-in tools to ensure quick and easy diagnosis and repair.

## N+1 Redundancy and No Single Point of Failure

Look for a solution with a distributed architecture that allows for the use of n+1 redundancy, which means that extra parts—as opposed to entire units—can be added to provide redundancy. Some vendors have 1:1 redundancy, which means twice the hardware is used to accomplish redundancy. Other systems use n+1 redundancy—which improves reliability since it is not doubling the hardware. For instance, the n+1 redundancy solution may need two extra units (where parts to the IP telephony system are duplicated within the two units), while a 1:1 redundancy solution needs five extra units because each unit is duplicated in its entirety. Essentially, using n+1 redundancy creates a multi-unit system with no single point of failure.

In addition to a distributed architecture that provides n+1 redundancy, look for a solution that interconnects each module using IP rather than cards in box slots. This design uses the Internet as a bus rather than having a proprietary backplane, which allows you to use a wide variety of chips and software and also reduces the costs and increases speed because of the use of IP and Ethernet. This design also allows you to seamlessly scale your system to meet organizational growth demands, just as the Internet allows for growth. Finally, look for a system that provides most of its feature upgrades via software so that there is minimal time between the release and your organization's use of these features.

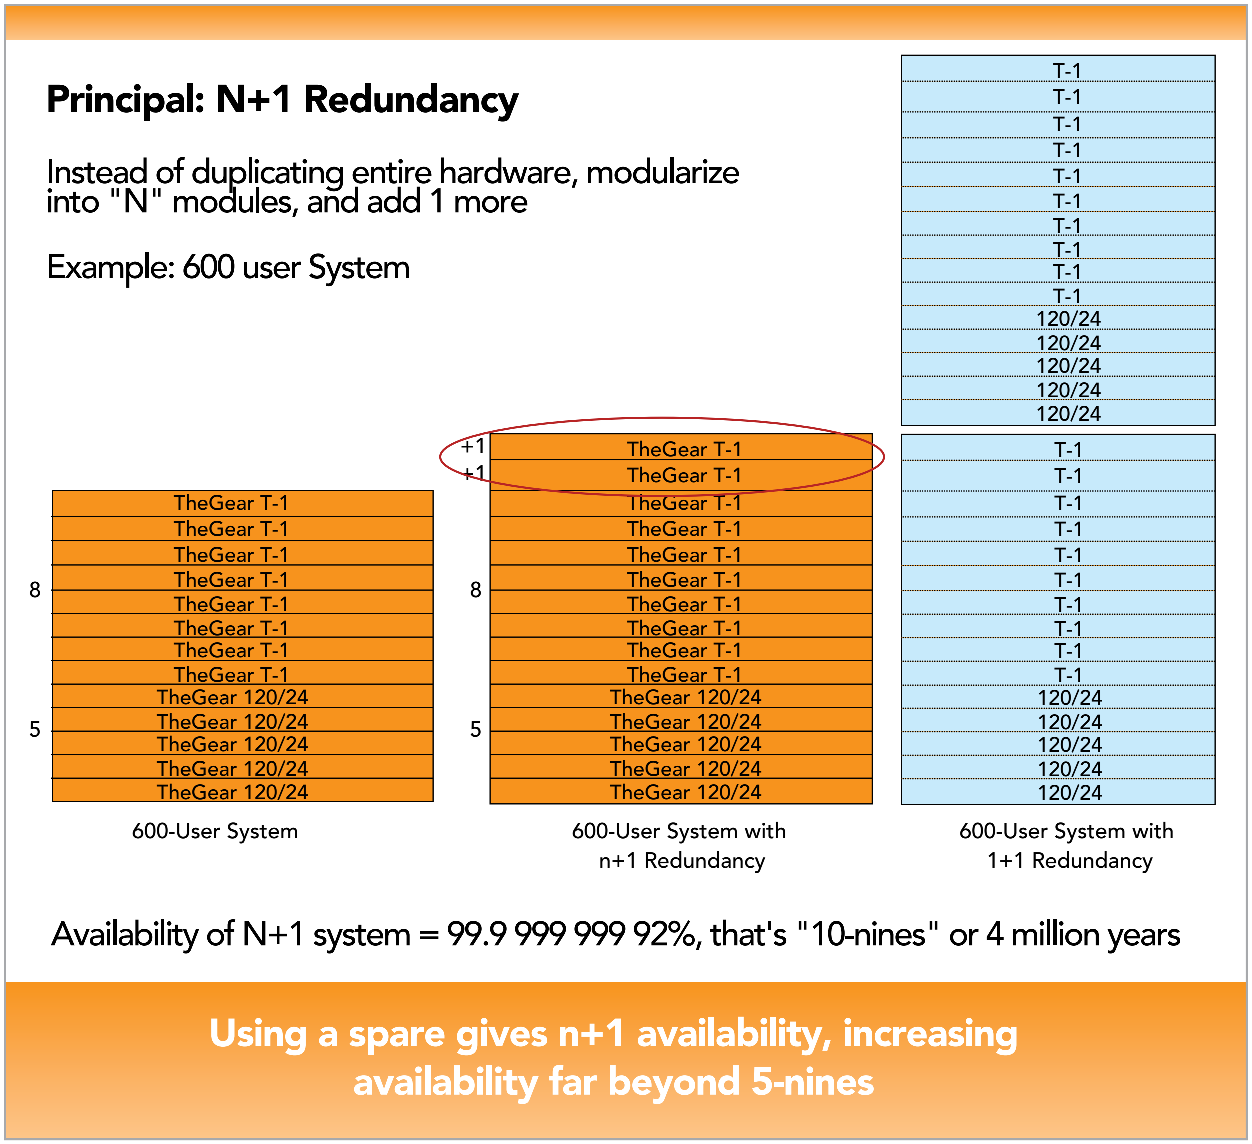


The goal of five-nines reliability is impossible for most systems because redundancy requirements can be complex and expensive. Using n+1 redundancy is not only more cost-effective, but it is less complex, which in turn reduces the chance of failure.

## Network Reliability

The biggest hurdle when implementing an IP telephony solution is ensuring it works properly with the existing underlying infrastructure. LANs and WANs have lower reliability than telecommunications systems and are prone to quality-of-service (QoS) issues that make VoIP solutions unreliable. LANs have multiple serial components, which negatively affects the reliability (typical LANs achieve three to four nines of availability), but it is possible to achieve five-nines availability on a network by using a redundant aggregation switch with redundant paths. After all, four-nines reliability translates to two hours of downtime per year. Can your organization afford that? Most 24/7 operations cannot. Focus on solutions that allow these redundant paths to an aggregation switch.

WANs cause the biggest headache because WAN links are generally available only 99% to 99.9% of the time, and voice quality availability can be as low as 98%. If your employees depend on superior voice quality for their many conference calls, for example, this is going to be a problem. Some solutions exist that distribute call control to local switches, which means that if a WAN link goes down, a remote switch can handle the calls because call control, business logic and system database information are all available within that switch.

A system with centralized call control relies heavily on its WAN connection because when it goes down, remote sites have no call control, which means calls cannot be made unless a backup system is in place. Look for a distributed solution that provides full and seamless call control functionality even during a WAN failure.

## Application Reliability

In addition to ensuring your system is reliable in terms of hardware, you must also ensure that VoIP system applications, including auto-attendants, voice mail, and desktop integration, work all the time for your employees. Look at systems that offer one application server for a full range of applications. You can use more than one server depending on your organization size, but make sure that it is not one feature per server, like some solutions may force you to do. A truly reliable system, in terms of applications, uses a site hierarchy, which means the first application in a user's hierarchy is used, and each application server has access to the configuration database in a central server. This design is highly reliable because each application server caches the configuration database, making information and applications available even during network downtime. For example, in the case of a network outage, remote users with their own server are unaffected by a failure in another server so that individual sites can serve features like auto-attendant.

## The Bottom Line

There is always the possibility that a system can be completely unreachable because of multiple LAN and WAN problems (remember the saying, "never say never"). Look for solutions that allow you to build into your system a backup plan, such as the ability to implement failover trunks, switch failover, and copper bypass for emergency service. There are lots of vendors out there offering piecemeal solutions that could leave you dealing with increased complexity and decreased reliability. A distributed architecture is a good fit for multi-site organizations, and n+1 redundancy designs will keep your costs—and your chances of failure—way down. The next chapter will go over system handsets, including analog and IP telephones as well as hard and "soft" phones.