IP Strives To Match TDM's Reliability

David Yedwab and Chris Talbott

We won't have converged networks until operators are confident IP can deliver. The case is being tested even as we speak.

Since the 1990s, the idea that converged networks and IP could ultimately replace TDM networks for voice as well as data has been toyed with and then dismissed. The basic reason was simple: Data networks—IP networks in particular—were thought to be too unreliable to meet the expectations of users accustomed to traditional TDM voice networks.

We take the availability and reliability of voice networks for granted, and feel quite comfortable routinely betting our lives on them. Even when power goes out in a storm, the phone almost always continues to work. (Our friends who work for our local power company say the reason is because the power lines are always highest on the pole, so the trees fall on them first.)

The phrase that characterizes the availability of digital TDM networks is "five-nines"—99.999 percent availability. How far are IP networks along the development curve toward that goal? The answer may surprise you: Within 18–24 months, there'll be parity in availability for well-designed and managed IP networks.

Based on our investigations for this article, this parity will be made possible through improvements brought about by Multiprotocol Label Switching (MPLS), extended features of protocols such as IS-IS and OSPF and emerging hardware/software combinations for rapid fail-over. Together with a growing base of real-world experience in managing IP networks, the pieces to solve the converged networks puzzle are falling into place.

The path to building high-reliability IP networks is clear, although many steps will have to be taken in parallel to improve the availability of IP networks. We need to run the best IP technologies in production environments, which, paradoxically, requires the PSTN for rapid fail-over. As an interim step, the parallel-networks approach offers the advantage of resilience through diversity, which is a precursor to the construction of a converged network with the survivability inherent in the Internet concept—i.e., the ability to route communications around damage to the network, as we saw after 9/11.

Definitions

When we talk about high-availability or "carrier-grade" IP networks, what do we mean? The generally-accepted definition is 99.999 percent or greater—around 5.3 minutes' downtime per year.

It's important to note that standard definitions of availability differ between TDM and IP in ways that affect a fair comparison. "All circuits busy," for example, is not considered an outage on a TDM network, but it would be on an IP network.

We also need to point out that we're talking about networks controlled and managed by, if not a single service provider or enterprise, then a small, finite number of entities. We're not talking about the broad, unmanaged Internet. Even multi-vendor/multi-network situations still provide more control than with the public Internet.

It's also helpful to understand the basic equation used for all factors that affect availability. The overall availability of the network is, in general terms, the product of this equation for each element:

$$\%\ \text{Availability} = \frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}} \times 100\%$$

By determining the total time a network element is in operation (Mean Time Between Failure, MTBF), and dividing by the sum of time in operation and Mean Time To Repair (MTTR, or the time the element is not working), you can determine the portion of time the element is available—how many "9s" it has. You also can substitute "customer-impacted minutes" and "total available minutes" in this equation. By determining the product of the availability for all the rele-

David Yedwab is executive vice president of The Eastern Management Group. Chris Talbott is a director with The Eastern Management Group.

vant elements, you can determine overall network availability. (For more detail on this formula see *BCR*, May 2002, pp. 22–27.)

The Myth And Its Origins

Computers are endlessly patient. People are not. Computers have whatever expectations we tell them to have; human beings have large and constantly growing expectations. That sums up the issues that arise as we converge data and voice networks.

Factors like jitter and latency are not fatal to the exchange of data, but human expectations of a voice conversation are different. Delays of more than about 50 milliseconds create annoyance for human brains trying to understand what someone is saying.

Even as digital technology began to enter voice networks in the 1970s, the engineers knew that there were human factors they had to accommodate. While the efficiencies of transporting voice in digital packets were becoming increasingly compelling, the new technologies were deployed to recreate the virtual equivalent of an old-fashioned, end-to-end copper connection, a circuit that stayed nailed up for the life of the conversation. The fact that we used Time Division Multiplexing (TDM) to sneak bits of other conversations into long (from a computer's point of view) silences didn't alter that.

The early gains in availability in digital TDM networks resulted from such things as improving multiplexing technologies of the copper network. As fiber entered the network, moving from multimode to single-mode, quality improved again. In the 1980s, moving from asynchronous technologies to SONET made its contribution, and improvements to SONET raised the bar yet again. With "survivability" assured, the path on which a conversation was traveling could be lost without losing connectivity; traffic was automatically switched to protection circuits so quickly that callers never missed a syllable.

But digital TDM and SONET didn't deliver five-nines from Day One; the SONET software introduced in the late 1980s didn't reach stability for two or three years, and it took roughly 15 years —until the mid-1980s—for digital TDM networks to reach five-nines of availability. Relative to reliability, IP also will go through a maturation process.

IP Networks: Where We Are today

Beginning with the drive to develop softswitches in the mid-to-late 1990s, telecom engineers started to seriously examine how to build an IP-based network that could carry voice traffic with the same level of quality that users of voice networks have come to expect. It's important to note that mainstream service providers have been battle-testing a complete set of pure IP technologies for only a few years.

But instead of the 15 or so years it took to push digital TDM into the five-nines territory, IP maturation is expected to achieve the goal in roughly five years, and we're already past the halfway point. Incremental improvements in hardware, software, processes and network management will continue to reduce cost and complexity, and those working on these issues say the path to these improvements seems clear. Indeed, in sufficiently controlled situations—e.g., single-vendor networks—IP networks can achieve five-nines today.

In Figure 1, we overlap the two maturation curves. This illustrates, qualitatively, the median between the most and least conservative views we hear. The gap between today's IP and TDM networks has been bridged in theory, in the labs and in alpha and beta tests. What's missing is real-world experience to soothe the healthy skepticism of service-provider engineers and managers. However, that experience is accumulating.

Analyzing and improving IP network availability begins by creating specific categories: hardware, software, environmental considerations, network design and, last but certainly not least, the combination of process design and human error.

- **Hardware:** Hardware is, relatively speaking, the easiest factor to predict, and therefore the easiest to fix, given appropriate resources. From the smallest component in each device on up the hierarchy, we compute the availability of each piece of the network making up a particular data path, and combine them to show overall network availability. Equipment providers have developed spreadsheets and other software tools to make this job easier, and they understand the common pitfalls. We'll give a summary of the major factors here.

First, you must break the network down into scenarios, based on the paths that data will take through the network for your chosen users. By creating block diagrams showing all the network

It took roughly 15 years for digital TDM to reach five-nines availablity. It will take IP about five years.

FIGURE 1 Maturation Curves

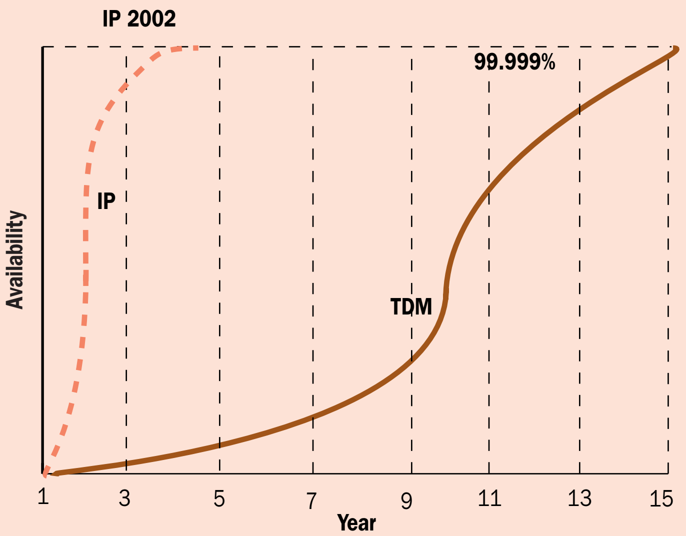

Boosting IP Availability Now

Given the potential leverage of small improvements, make sure you examine the most common ways to boost network availability:

■ Measure comparative MTTR rates for the device types in your network. Some hardware may need upgrading, while you may want to revert to earlier, more stable versions of some software.

■ Consider upgrading service contracts to ensure faster delivery of the right parts and expertise, particularly those that can potentially affect a large portion of the network.

■ Improve processes for handling network failures; review historical data on your fastest and slowest restorals for a given type of failure to identify targets for improvement; process-induced human errors are frequently a major contributor to long outages.

■ Examine appropriateness of fail-over mechanisms. For example, OSPF works fine for data-only environments, but will not be fast enough for the demands of voice traffic on a converged network.

■ Look for root causes of failures; for example, does a process repeatedly create the same human error? Requiring manual entry of configuration information is virtually guaranteed to create problems.

■ Ensure that change-management processes are well thought out and reflect historical experience with failures. A very large percentage of IP-network failures occur during changes—software or hardware—usually because those involved do not have the right information or do not take adequate safeguards—e.g., rolling back to stable software if an upgrade fails□

elements for these paths, you can see where devices are in parallel and serial arrangements (for more on the impact of parallel vs. serial configurations, see *BCR*, May 2002, pp. 22–27). You can then create a list of device types, so that you can develop availability estimates for each.

From there, it is a straightforward, albeit laborious task to combine availability for individual network devices, including groups of parallel devices converted into serial availabilities. The product of all these individual availabilities yields the overall network availability.

As you go through these calculations you find a kind of "butterfly effect," as they say in chaos theory: Small factors—a single faulty router card or badly designed restoral process—can have very large leverage on overall availability. This is exquisitely double-edged, of course. It means that eliminating one hardware weak spot in a network or one faulty process design can jump a network from four-nines to five-nines all by itself. It also means one unanticipated or unnoticed component failure can drop a five-hour outage onto your network.

The good news is that we know how to get MTBF and MTTR data, whether from device manufacturers or from experience with our own networks. Given that, hardware availability problems are eminently solvable.

Finally, we note that the biggest mistake made concerning the hardware factor is assuming that it is the only important factor in IP network availability. It's not; indeed, a thorough availability analysis only begins with hardware. Valarie Gilbert, who heads Cisco's U.S. customer program management office, put it this way: "Our data shows that hardware is only about 40 percent of the problem. Eliminating the largest contributions to network downtime means looking closely at recovery procedures, network design, sparing, software, etc."

■ **Software:** In an IP network, there are two major types of software: the operating systems (OSs) in individual devices such as routers, and the network control software. Both are often run in parallel designs, greatly improving their contribution to network availability.

A common example of an IP network failure would be for the OS in a router to crash. We can find out how quickly the router can restore itself after a crash, rebuild its routing tables and start passing data again: Mean Time To Repair. We can thus calculate the availability of each device in terms of software, adjust network design as required—including parallel design, dual homing for switches, and so on—and thereby mitigate the effects of software failure on overall network availability.

The takeaway here is that reliable software running in redundancy does not add significantly to the downtime of the network, since the parallel element can take over, assuming the fail-over mechanism is well designed. The key word, obviously, is *reliable*. The experience of others, as well as what we can learn from running the software in a test situation, can help. This also reinforces the desirability of being able to get back to a reliable version quickly if necessary after an upgrade.

There is no reason that software for IP networks should be any less reliable than software for TDM networks, assuming the level of testing is the same, as it should be. IP networks also can benefit from the experience of TDM network managers very directly: The most important way to avoid software-driven network failures is to create and aggressively manage upgrade procedures.

In a nutshell, upgrades should happen when network traffic is lowest, the most skilled personnel available should conduct them and the process should include the ability to instantly restore the previous software in the event of a problem.

Environmental considerations: These don't require much discussion, for one simple reason: They are virtually the same considerations you would have for a TDM network, for both voice switches and instruments. Although early IP voice equipment was not line-powered, the current IP telephony products are all line-powered. As with traditional networks, then, the power issue focuses on the equivalent of the central office, where the same sets of processes and tools can be employed: battery backup, generators, etc.

The estimation of needs follows the same well-known path from obtaining accurate power-availability data from utilities, based on geography and other site-specific factors, to providing the appropriate balance of processes and tools to bring this factor up to five-nines. Since the methodologies and historical data for this are well known, we will say no more here.

Network design: With good data in hand about current network availability, we can add parallel elements to bring weak points up to five-nines. The caveat is that parallel designs require fail-over mechanisms, which, per Murphy's Law, can also fail. Altogether, however, experience with both IP and TDM has shown that parallel network designs—monitored, periodically measured and adjusted—can be a powerful way to achieve five-nines in production networks.

The elements that should be deployed in parallel depend on factors specific to the individual network and the management processes employed. Individual routers may have processors running in parallel, or two devices with individual processors may be used. The layout of facilities connecting network nodes will vary, thus changing the level of guaranteed availability, depending on how backup facilities are engineered.

While parallel designs increase the initial cost of the network, they also greatly increase availability. One popular rule of thumb: Getting from four-nines to five-nines doubles your costs, but gives you 10 times the availability. The key to maximizing this tradeoff, of course, is to have the historical data to show which devices in the network will yield the greatest measurable benefit from parallel design in terms of overall network availability.

What trends does the data show at the moment? Two factors prevent that story from being told right now. First, the equipment vendors and service providers are still writing the first draft of this history. Second, what those companies are learning is valuable information that they don't want to share with competitors. As demand rises for converged networks, and common practices proliferate, rules of thumb will no doubt emerge.

FIGURE 2 Analyzing And Human Errors

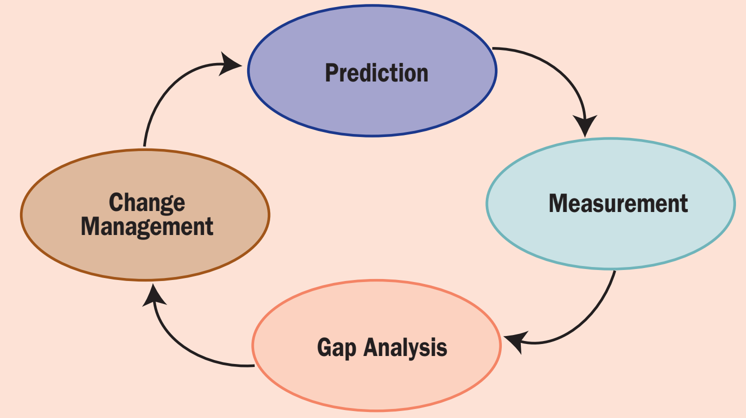

Processes and human error: These two subjects should be talked about together since they are inextricably linked. Two key process areas to examine for best practices are, first, the combination of software change management and software version uniformity, and second, configuration consistency, across common platforms in the network.

"We get customers started by addressing the low-hanging fruit around recovery and operational procedures," Cisco's Gilbert said. "By themselves, these can nearly double network availability. Once we've captured these quick wins, we….[drive] specific design, software and management features that improve availability."

The steps to minimize human error are the standard ones for improving any process—prediction, measurement, gap analysis and change/change management. Since this is an iterative process, they should actually be shown as depicted in Figure 2.

Step 1—Prediction: Begin by predicting what the availability of your network should be, all other factors being equal. This is the entire process outlined above, and includes metrics for hardware, software, environmental factors and network design.

Step 2—Measurement: Track, over time, the causes of all failures, including those caused by process/human error, as well as details on total operating time, number of failures and network downtime. This data can then be analyzed and compared in Step 3 with the baseline (established in Step 1).

Step 3— Gap Analysis: Compare the predictions from Step 1 versus the reality measured in Step 2. Root causes—usually patterns that depict process problems—should become apparent, or at least the places to start looking for them should.

Step 4—Change/Change Management: Create methods for solving the problems in a way that doesn't create new problems at the same time. If a software upgrade is in order, for example, make sure a well-honed rollback process is already in place. Create and follow a set schedule for the improvement cycle.

What's Next

Much work is being done by the Internet Engineering Task Force (IETF) and elsewhere to bring IP protocols to a higher level of maturity. Sources at major service providers believe the protocols will reach the required maturity within 24 months.

One reason they're optimistic is MPLS, which, allows the kinds of traffic management for IP networks that TDM networks have evolved over the years. MPLS sets up specific paths for given strings of packets by labeling them so routers can forward them more quickly. In the OSI model, this means forwarding most packets at Layer 2 rather than Layer 3. It also allows for QOS, setting priorities for different types of traffic—delay-sensitive voice ahead of email, for example.

However, MPLS options also address network restoral. Fast Reroute, LSP (Label Switched Path) and Backup Record Route each rely on different methods to provide the kind of rapid recovery that ATM and SONET have today. Each has pros and cons, and all are currently in the throes of industry standards evolution—either through official bodies such as IETF and ITU, or through the actions of equipment makers who are responding to market pressures.

In parallel with the work on the IP standard itself, there's work underway to develop new technology to make IP networks more reliable. These development efforts include:

- Additional hardware and software in the IP network that automatically switches to alternate routes in the case of facilities failure. This must happen quickly enough that it is not noticeable to voice callers, as TDM networks routinely do today. All the major router and switch vendors are currently working on this objective.
- Routers, switches and other hardware that can be software upgraded without being taken out of service for 15 to 45 minutes, as is the case today.

These planned improvements to router technology aim to give IP devices the ability to switch to backup equipment transparently to the data flow and therefore to end users. For example, individual router line cards will be able to reboot without taking down the rest of the gear in the same box with them. Switching to redundant route processors in a "stateful" way will make possible a fail over, handled at Layer 3, that is smooth enough not to interrupt a voice conversation.

While we are concentrating here on the IP layer, it also should be noted that similar progress is being made at the physical layer, in next-generation optical technologies. Many optical components today operate at well above five-nines. As with IP, new mesh architectures and switch designs are already in use in some production environments.

Together with upcoming versions of current protocols, these provide the rapid (< 50 millisecond) switch to alternate facilities required for voice applications, and when combined with the progress in the IP layer described above, give IP the potential to leapfrog TDM since there will be no need to over-engineer the network to provide fully redundant facilities (as in SONET rings, for example). Examples of the protocols being improved to meet these needs would be Border Gateway Protocol (BGP), Hot Standby Routing Protocol (HSRP), Intermediate System to Intermediate System (IS-IS), Open Shortest Path First (OSPF) and Resource Reservation Protocol (RSVP).

Conclusion

We're in a classic chicken-and-egg dilemma in the evolution to high-availability, converged, multiservice IP networks: Without assurances that production IP networks will work as well as TDM networks, service providers and enterprises are reluctant to rely on them. Yet without those production environments and the real-world data they provide, progress remains slow.

Fortunately, however, the industry is taking steps out of this quagmire. ATM and frame relay networks, analogous but not identical to pure IP, have not only met but exceeded the availability of digital TDM in production environments. Some well-publicized failures provided the incentive to understand the real-world problems involved, especially the role of processes and human error.

Moreover, leading service providers are deploying IP equipment for converged networks, although they don't publicize these projects very much. This is understandable given the early stages of the network evolution, and because they all have individual ways of overcoming the transitional issues in converged networking.

The ultimate takeaway here is simple: The theoretical understanding, the hardware, the software, the processes and even most of the real-world lessons required to achieve IP networks with five-nines availability exist today. What remains to be done falls into three broad categories:

- Experience with the coming mature versions of MPLS and related protocols.
- Development of the emerging crop of hardware/software combinations with rapid fail-over capabilities.
- Development of new, appropriate management processes□

Company Mentioned In This Article

Cisco (www.cisco.com)