

Beyond Five Nines: The New Age of “Uber Reliability”


With the advent of the age of “uber reliability,” best-effort services are passé. Even the long-vaunted “five nines” gold standard of reliability may no longer be good enough for some end users and applications.

Executive Summary

The benchmark for telecom reliability is 99.999% (“five nines”), which is equal to about five-and-a-half minutes of downtime per year. The five-nines benchmark is a common way to quantify “carrier class” — one of the most widely used — and misused — adjectives in the telecom industry. With no real established measurement standards, vendors can claim products are carrier-class without having to detail under what operating conditions and circumstances the products were measured. That makes the importance of standardizing on “real world” performance data critical, giving service providers, government agencies and enterprises a common tool to perform reality checks and compare telecom infrastructure product reliability.

Performance figures are important because they directly impact a service provider’s or end user’s bottom line in a variety of ways. For example, chronic equipment failures can increase a service provider’s customer acquisition and retention costs, while a service provider with highly reliable infrastructure can use its service quality as a powerful market differentiator. Although reliability has always been a factor in a service provider’s competitive position, it has become an even larger one as the market for “best effort services” dwindles among both consumers and enterprises. Consumers and enterprises now expect better than best effort services. A February 2007 survey of 22,000 European mobile users conducted by M:Metrics, now part of comScore, an independent market research firm, found that unreliable service is a major reason why early adopters of mobile TV and video have dropped those services.

Meanwhile, many regulators have spent the past few years increasing their reliability requirements for service providers, largely as a reflection of the fact that so many consumers and enterprises depend on telecom services. One example is the U.S. Federal Communications Commission (FCC), which now requires wireless carriers to report Mobile Switching Center (MSC) outages that last longer than 30 minutes, regardless of the number of customers affected. In some cases, the FCC now requires service providers to report outages that affect no customers, such as a circuit that loses its backup.



This white paper discusses why telecom is entering a new age of “uber reliability,” in which best effort services are an increasingly tough sell and five nines may no longer be good enough for some end users and applications. This paper also looks at how the telecom industry can define and measure reliability in ways that enable the accurate comparisons necessary to make informed decisions about which infrastructure products are the most reliable. Finally, the paper provides an overview of Tellabs reliability initiatives and results, such as the five nines plus reliability of the Tellabs® 7100 Optical Transport System (OTS) and the Tellabs 5500 Digital Cross-Connect System (DCS).

The importance of standardizing on “real-world” performance data is critical in order to give service providers, government agencies and enterprises a common tool to perform reality checks and compare telecom infrastructure product reliability.

¹ www.mmetrics.com

“Uber Reliability”

Although “carrier class” may be one of the most widely used adjectives in telecom, it also has potential to be misused, to the point that any vendor can claim that its product is carrier class without having to prove that assertion. But savvy service providers understand the value of real-world performance data — value that can be measured in terms of impact on the service provider’s bottom line.

In addition to direct costs, chronic network failures and lack of reliability can have an impact on service providers’ investment prospects. Investors closely monitor churn and customer acquisition costs, so a service provider with unacceptable churn rates will have reduced access to capital, further undermining its competitive position.

Although reliability has always been important for service providers, two trends are steadily increasing its value even further. The first is the evolution of telecom services — including wireless and Web-based communications — from a luxury to a necessity among both enterprises and consumers. That evolution has winnowed the market for “best effort services” — shorthand for a lesser grade of quality than carrier class.

The second trend is increased competition, with more service providers vying for a finite pool of potential customers and revenue streams. Even in developing countries, consumers and enterprises typically have multiple choices for telecom services. A service provider’s reputation for reliability — or lack thereof — is one factor that consumers and enterprises in developed and developing markets consider when choosing a service provider.

This analysis comes through in a February 2007² survey of 22,000 European mobile users by M:Metrics. The study found that unreliable service is a major reason why early adopters of mobile TV and video have dropped those services. An average of nearly 24% of all respondents said that concerns about service quality and reliability were the main reason why they discontinued service. In some countries, the percentage was even higher, such as 29% of U.K. respondents.

This experience is causing a ripple effect by creating a wider marketplace perception that mobile TV and video are generally unreliable. For example, among the survey’s U.K. respondents, 6% had never used mobile video and TV, yet they cited quality and reliability as reasons not to try those services.



For mobile operators and their content provider partners, these findings are both a problem and an opportunity: They’re a problem because the wide perception of unreliability undermines the market for every operator’s mobile TV and video services, which are an important way to drive new revenue and offset flat or declining voice revenue. But the negative findings also indicate an opportunity to create mobile TV and video services that are so reliable that they stand out simply for that reason, thereby attracting more users.

Although reliability has always been important, two trends are steadily increasing its value: the evolution of telecom services from a luxury to a necessity among both enterprises and consumers, and the increased competition of more service providers vying for a finite pool of potential customers.

Simply reducing the price of mobile TV and video services may not be enough to win over both people who have tried those services and those who have never used them. “Pricing has already been highlighted as a stumbling block for recurrent use of mobile video and TV services, but we were surprised by just how much value users place on quality and reliability,” says Paul Goode, senior analyst at M:Metrics. “Once the basic requirements of quality and reliability are good enough, the focus will rightly shift to issues of programming, brands and marketing in addition to price. This research highlights the need to address quality and reliability so the industry can retain viewers, which is a key part of growing audience numbers.”

² www.tellabs.com/news/2007/nr021207.shtml



Mobile commerce is another example of the importance of reliability. Unless financial institutions and their customers are confident that transactions won't be lost or compromised because of network problems, mCommerce won't live up to its potential: \$22 billion worth of transactions by 2011, up from about \$2 billion in 2007, according to Jupiter Research.³ Wireless carriers and their investors are counting on mCommerce to drive new revenue, but that opportunity is available only to operators capable of providing the Service Level Agreements (SLA) and other guarantees that financial institutions demand from their wireless partners.

Regardless of the type of application or service, reliability is important to all consumers. If it weren't, service providers such as Verizon Wireless wouldn't be able to successfully play up the reliability of their networks as a market differentiator. Younger demographics such as Generation Y and Echo Boomers have shown a willingness to churn when their wireless or Internet service is unreliable, and they've also shown a willingness to pay a premium for reliable service. A 2007 Forrester Research, Inc. survey showed that 23.6% of EchoBoomers (18- to 24-year-olds) and 22.3% of GenY customers (25- to 29-year-olds) would switch to a new Internet Service Provider (ISP) if reliability is poor with their current provider.⁴ As one gamer put it, he's "willing to pay even double to get more than double the speed and reliability. If it was actually a gaming package that guaranteed better pings with considerably better speed, I would pay up to double ... If the speed and reliability fluctuated with user loads, I would not use that ISP."⁵ That's one example of how service providers can use reliability not just as a market differentiator, but also as a way to drive additional revenue and justify a price premium in a commoditized market.

³ www.tellabs.com/news/reprints/emerge_winter07_money_reprint.pdf

In the last several years, as consumers and enterprises have come to depend on telecom services to a greater extent than ever before, regulators have increased service provider reliability requirements. Since January 2005, the FCC has required wireless carriers to report any MSC outages longer than 30 minutes, regardless of how few customers are affected. Cable telephony providers also have a new FCC mandate to report outages that impact more than 900,000 total user minutes. Under that guideline, a cable telephony outage lasting just 15 minutes would fall under the FCC outage reporting rules if it affects more than 60,000 customers — a likely scenario considering the size of some cable operators' telephony customer bases. Now there are even some cases in which the *FCC requires service providers to report outages that affect no customers*. For example, a circuit losing its backup must be reported, even with zero direct impact on customer services.

Reliability is now a regulatory requirement, a marketplace requirement and a key market differentiator — not just for service providers, but for infrastructure suppliers too.

The result is an age of “uber reliability,” in which reliability is both a regulatory requirement and a marketplace requirement. In the process, reliability has become a key market differentiator — not just for service providers, but for their infrastructure suppliers, too. Reliability also should be viewed as one major part of the Total Cost of Ownership (TCO) for a piece of network equipment because of the way that chronic problems can increase customer acquisition and retention costs.

Defining and Measuring Reliability

Historically, “five nines” or 99.999% uptime has been considered the ultimate in availability. Five nines translates to approximately just five minutes and 35 seconds of total system downtime in a calendar year which, carried across a large customer base, could mean an outage of just a few seconds affecting a major portion of the customer base in an event that degrades system reliability.

A variety of factors can lead to unreliability. One common cause is untested code paths, which can occur when a service provider's test bed doesn't cover every possible scenario that will occur in a real world deployment. Another is when a network element is taxed to the point of failure by traffic loads that it was never intended to handle.

“With millions of paying customers relying on them, failure isn't an option for carrier networks,” says John Mazur, principal analyst of switching and routing at Ovum, an independent analyst firm. “Public networks have the most demanding

⁴ “North American Consumer Tech Online,” Forrester Research, Inc., 2007.

⁵ <http://www.inspirethenewlife.com/people/synide.cfm>

requirements for both scalability and reliability, but can only be as reliable as their weakest network element. Network operators must rigorously test all network elements to understand where the weakest link is in their network, which could cause cascading failures. The migration of IP switching and routing technology from enterprise LANs to carrier-class critical public network infrastructure has not been a small undertaking.”

That analysis highlights the value of working with an infrastructure vendor that has a reputation for reliable products. Service providers should also look for vendors that design their products to mitigate problems created by the fact that there are so many variables in the real world. A prime example is features that enable fast recovery from a failure. “Carriers must serve all customers and can’t put restrictions on their traffic, thus are at risk of failure due to unforeseen customer traffic characteristics. Network elements features such as non-stop routing and fail-safe operation based on redundant processing enhance the system reliability of public networks,” Mazur says. Service mirroring is supported on all Tellabs® 8800 Multiservice Router (MSR) Series platforms, including the Tellabs® 8860 MSR, Tellabs® 8840 MSR and Tellabs® 8830 MSR.

The goal of the Tellabs “reliability by design” philosophy is to develop products to five nines standards for worst-case scenarios. In more typical scenarios, many Tellabs products may deliver eight nines reliability — equal to just .32 second of downtime in a year.

One challenge that service providers face when trying to assess reliability is that it’s difficult to collect the information necessary to make an apples-to-apples comparison between multiple vendors’ products. For

example, different vendors often use different metrics when measuring outages and reliability. The term “outage” also can have different definitions. For instance, the FCC requires Automated Reporting Management Information System (ARMIS) reports of telecommunications switching systems to count only outages longer than 10 seconds. That’s because failures that last less than 10 seconds often are invisible to end users. So one viable way to assess reliability is first to determine the length of time before a failure becomes noticeable to end users.

The next step is to identify a recommended parameter for equipment reliability. This measures the probability that the equipment will not fail at any time during a certain period. This figure can be used to set the requirement, such as the traditional telecom benchmark of 99.999% uptime (five nines).

The network element then can be tested to determine whether it can meet the requirement. There are several options for these tests:

- Collect field data on similar products (or earlier versions of the product). Calculate the proportion of times the equipment fails during a certain period and use that collected data to calculate the confidence interval (the range of values in a statistical study that contain the true parameter value within a given probability).
- Run an accelerated life test of the equipment. Use an Arrhenius model with activation energy equal to 0.4. The 90% confidence interval for the failure rate should include 20,000 Failures-in-Time (FIT) and should not exceed 60,000 FIT.
- Perform a reliability prediction using Mild-Std-217F or Telcordia SR-332 standards. The reliability predictions should have a failure rate less than 20,000 FIT.

Tellabs — Reliability by Design

Although the service provider market has long used five nines as a benchmark, other markets have begun adopting this requirement. One example is the U.S. government. According to a recent FedSources study, *Changing Government Telecom Network Requirements for the 21st Century*,⁶ one key trend is that federal networks now must offer five nines reliability. In fact, the U.S. General Services Administration’s (GSA) new Networkx⁷ contract, awarded in 2007, increases the baseline requirement to five nines from 99.8% reliability in its predecessor, FTS 2001. This difference can be expressed in time: FTS 2001 allowed more than 17 hours of downtime per year, while Networkx allows only five minutes.

The federal government’s increased requirements highlight a key difference between enterprise-grade and true carrier class network equipment. As FedSources Senior Vice President and Chief Knowledge Officer Ray Bjorklund noted, the increased requirements “challenge the prevailing belief that government telecom customers should ‘think and act’ like enterprises. By the end of the [study], the evidence clearly showed that many government telecom networks have a scale and complexity that cannot be adequately addressed by enterprise-grade equipment.”

Another recent study echoes this finding. A February 2007 survey⁸ of 388 federal government IT managers found that 93% require 99.999% telecom network uptime. Conducted by PostNewsweek Tech Media on behalf of Tellabs, the survey also found that reliability and outages are among the main issues that keep them up at night. As one respondent put it: “Economy and reliability are equally important. In some cases, reliability is more important.”

⁶ <http://www.tellabs.com/news/2007/nr010907.shtml>

⁷ www.gsa.gov/Portal/gsa/ep/contentView.do?contentType=GSA_OVERVIEW&contentId=16100

⁸ www.tellabs.com/news/2007/nr022607.shtml



One recent example of Tellabs' ability to meet the U.S. government's high reliability standards is the Tellabs® 7100 Optical Transport System (OTS). In December 2007, the United States Information Systems Engineering Command's (ISEC) Technology Integration Center (TIC) certified the Tellabs® 7100 OTS for use in an environment that's the ultimate in mission-critical applications: broadband connections between Army bases and soldiers in the battlefield.

The previous U.S. federal government standard for network reliability allowed more than 17 hours of downtime per year. The Network government contract awarded in 2007 specified a five nines requirement — 5 minutes, 35 seconds of downtime per year.

Another similarity between government agencies and service providers is that both typically require infrastructure that can deliver five nines reliability even in a network environment that has a mix of legacy protocols and

relatively new technologies such as IP/MPLS and Ethernet. That's because in both markets, the migration from legacy technologies to IP is a years- or decade-long process, with existing infrastructure investments that must be maximized before they can be retired. This situation highlights the importance of picking an infrastructure vendor that has experience not only with carrier-class networks, but that also with ensuring reliability amid a mix of new and legacy protocols.

Another key benchmark is the QuEST Forum's TL 9000 program,⁹ which provides service providers with an independent way to assess the quality of a participating vendor's products. The QuEST Forum's leadership includes major service providers such as AT&T, Bell Canada, Tata Teleservices, Verizon Communications and Vodafone, as well as Tellabs executives — an example of the value that industry sees in having an independently produced set of quality requirements.

The QuEST Forum is just one example of how reliability is part of the Tellabs DNA. To meet and exceed government, service provider and enterprise expectations, Tellabs has a philosophy of “reliability by design,” with a goal to develop Tellabs products to five nines standards for worst-case scenarios. As a result, in more typical scenarios many Tellabs products may deliver eight nines reliability — just .32 second of total downtime in a calendar year. Many vendors design products to achieve five nines reliability in ideal scenarios, leading to network failure when elements are taxed by traffic loads they were not intended to handle.

This design philosophy is one of the reasons why the worldwide service provider market perceives Tellabs as being highly focused on reliability. This perception comes across in an early 2007 survey of the eight largest wireline and wireless Tellabs customers in North America, as well as 24 of the company's largest service provider customers in the APAC, EMEA and Latin America regions. The majority of the respondents — 84% — work in network engineering and planning, and their feedback was based on their experiences with Tellabs infrastructure in the second half of 2006.

The survey's key findings include:

- When asked to rate overall product reliability vs. its closest competitors, Tellabs products rated an average of nearly 8 on a 9-point scale.
- 93% of respondents rated Tellabs overall product reliability as good or excellent. This response is an improvement over the previous year's survey, where the majority ranked Tellabs overall product reliability as good.
- Both wireline and wireless products ranked higher for overall reliability in 2006 than 2005.

These survey results show that Tellabs currently enjoys a reputation for providing highly reliable products. Nevertheless, what's exemplary today may be pedestrian tomorrow, as service provider expectations increase. That's why although service providers and many government agencies expect five-nines reliability today, many Tellabs products already meet what's likely to become their expectation in the near future: eight nines reliability.

⁹ <http://tl9000.org/index.html>

	Tellabs 7100 OTS	Tellabs 5500 DCS
Availability Minutes	15,601,902,086 wavelength minutes	57,617,200,000,000 DS1 minutes
Outage Minutes	3,350.18 wavelength minutes	6,883,515 DS1 minutes
Reliability Percentage	99.9999785271%	99.9999880530%

Table 1. Tellabs 7100 OTS and Tellabs 5500 DCS reliability, 2006–2007

Tellabs Performance Data — Real-World Reliability

To illustrate the representative reliability of the broad range of Tellabs products, we will examine reliability figures for both a newer Tellabs product and one with a long service history. The recently introduced product is the Tellabs 7100 OTS,¹⁰ which supports Add/Drop Multiplexer (ADM) and Wavelength Division Multiplexing (WDM) ring capabilities. Based on compiled outage data, the Tellabs 7100 OTS attained a 99.9999785271% reliability score over the two-year 2006–2007 span. Tellabs calculated these reliability figures by using wavelength minutes available and wavelength minutes out figures, all obtained using TL9000 reporting rules.¹¹

A more established Tellabs product, the Tellabs® 5500 Digital Cross-Connect (DCS)¹², which provides VT, DS1, STS-1 and STS-3c capabilities, has been widely used by major carriers worldwide for nearly two decades. The same customer outage statistics showed that the Tellabs 5500 DCS was 99.9999880530% reliable over the 2006–2007 period. Tellabs calculated these reliability figures by using DS1 minutes available and DS1 minutes out figures, all obtained using TL9000 reporting rules.

Both Tellabs products attained six nines of reliability, which translates to roughly 31.5 seconds of downtime per wavelength or DS1 per year.

Conclusion

With the advent of the age of “uber reliability,” best effort services are passé and even the long-vaunted “five nines” gold standard of reliability may no longer be good enough for some end users and applications.

An increasingly fickle customer base has shown the willingness to change service providers and even pay more to obtain premium quality and reliability, forcing service providers to demand the ultimate level of reliability when specifying and evaluating equipment. The reason is simple — to effectively compete and differentiate their services, service providers must be able to offer customers consistently reliable service. Their business boils down to the consistent availability of the equipment in their core networks. Tellabs is committed to continuing its 30-year history of designing equipment that’s consistently reliable in the field, in real-world conditions.

The Tellabs 5500 DCS and Tellabs 7100 OTS products attained better than six nines reliability in 2006–2007, translating to roughly 31.5 seconds of downtime per wavelength or DS1 per year.

¹⁰ www.tellabs.com/products/7000/tellabs7100.shtml

¹¹ These figures count only product-attributable outages. They do not include outages caused by the customer, external failures, or acts of God such as floods and hurricanes.

¹² www.tellabs.com/products/5000/tellabs5500.shtml

North America

Tellabs
One Tellabs Center
1415 West Diehl Road
Naperville, IL 60563
U.S.A.
+1 630 798 8800
Fax: +1 630 798 2000

Asia Pacific

Tellabs
3 Anson Road
#14–01 Springleaf Tower
Singapore 079909
Republic of Singapore
+65 6215 6411
Fax: +65 6215 6422

Europe, Middle East & Africa

Tellabs
Abbey Place
24–28 Easton Street
High Wycombe, Bucks
HP11 1NT
United Kingdom
+44 870 238 4700
Fax: +44 870 238 4851

Latin America & Caribbean

Tellabs
1401 N.W. 136th Avenue
Suite 202
Sunrise, FL 33323
U.S.A.
+1 954 839 2800
Fax: +1 954 839 2828

Statements herein may contain projections or other forward-looking statements regarding future events, products, features, technology and resulting commercial or technological benefits and advantages. These statements are for discussion purposes only, are subject to change and are not to be construed as instructions, product specifications, guarantees or warranties. Actual results may differ materially.

The following trademarks and service marks are owned by Tellabs Operations, Inc., or its affiliates in the United States and/or other countries: TELLABS®, TELLABS and T symbol®, and T symbol®.

Any other company or product names may be trademarks of their respective companies.

© 2009 Tellabs. All rights reserved.
74.2064E Rev. A 2/09