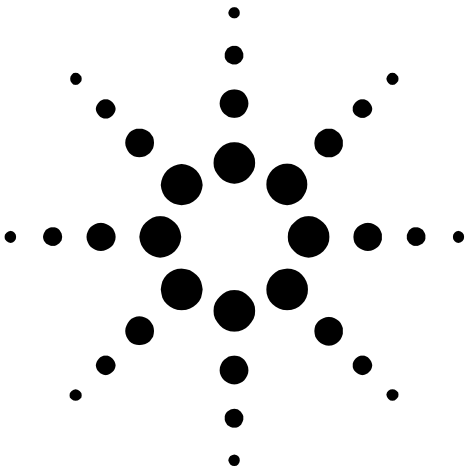

Technical Paper

Preventing IP Network Service Outages

Introduction

Service Providers would rather focus their attention on deploying new, revenue-generating services than on restoring network services. While some downtime is almost inevitable, “prevention is better than cure!”

This paper addresses the following issues:

- The real cost of network outages and service downtime
- Going beyond equipment vendor testing
- The seven stages of network technology adoption
- How testing each stage before deployment can reduce the risk of network outages

Agilent Technologies

The real cost of network outages

On April 13, 1998, in one of the mostly widely known and documented network outages, AT&T's Frame Relay network went down. Disruption to customers lasted for up to 26 hours and included the failure of mission-critical bank transactions. The fact that Service Level Agreements (SLAs) had promised 99.99% availability and four-hour recovery times did not go unnoticed. Approximately 6600 customers were not charged, resulting in millions of dollars in lost revenue alone.

AT&T shared its analysis of this outage with the FCC, the industry-wide Network Reliability Council, and other network providers. "By sharing this information and best practices," said AT&T's Chairman Michael Armstrong, "we will help customers avoid similar network outages no matter which carrier provides their Frame Relay service." (See reference [1] on page 10.)

The actual network outage examples described in this paper have been taken from published or publicly available sources. The intention here is not to criticize a group of Service Providers or manufacturers, but quite the opposite; those who have chosen to publicize the causes of outages should be applauded by customers for being open. They have shared their efforts to learn and improve the availability, robustness, scalability, and performance of network services. After all, every Service Provider encounters the same challenges.

Many of the costs to Service Providers caused by outages and downtime are tangible and easy to measure, such as

- Direct costs of service restoration and equipment upgrades
- Loss of revenue during outages
- Penalties associated with customer SLAs

However, the "real" costs include some losses that are harder to quantify but may be far greater. For example,

- Lost revenue from dissatisfied customers moving to competitors or taking new business to competitors
- The cost of a tarnished image; the lessened ability to credibly market future "premium" differentiated services and position them against competitors

According to USA Today, U.S. companies lost an estimated *$100 billion* due to network outages in 1999 alone. (See reference [2] on page 10.)

Beyond vendor testing

Before commercial release of a new network product or product upgrade, manufacturers put their equipment through a barrage of tests throughout development, system testing, and production. Some vendors go so far as to replicate parts of their customers' networks. However, Service Providers should be aware that vendor testing may not be adequate for their requirements, and need to consider the following questions:

- Will vendor products work bug-free in your unique network configuration?

  Manufacturers can't test every combination of features, topology, protocols, network size, ports, and services. This is an important part of evaluation and acceptance testing.

- Will they work in your multi-vendor environment?

  Manufacturers do not have the resources to test against every release of every manufacturer. Service Providers must test for interoperability with their own equipment and must also take into consideration equipment belonging to their own customers and to interconnect partners.

- Will they work with your legacy equipment?

  Manufacturers cannot be expected to test against every network element. Once again, interoperability with legacy equipment is a Service Provider responsibility.

- Market pressures push manufacturers to take shortcuts.

  Manufacturers do not release their detailed internal test procedures or test results. When caught in the vendor's rush to release a new device or capability, a Service Provider may find itself being the "guinea pig" for a new software release!

The seven stages of network technology adoption

Every Service Provider encounters similar challenges during the adoption of new network technology. To describe these challenges, it is helpful to divide the technology adoption lifecycle into the stages outlined in Figure 1. Below is a sample of the questions you should consider during each stage.

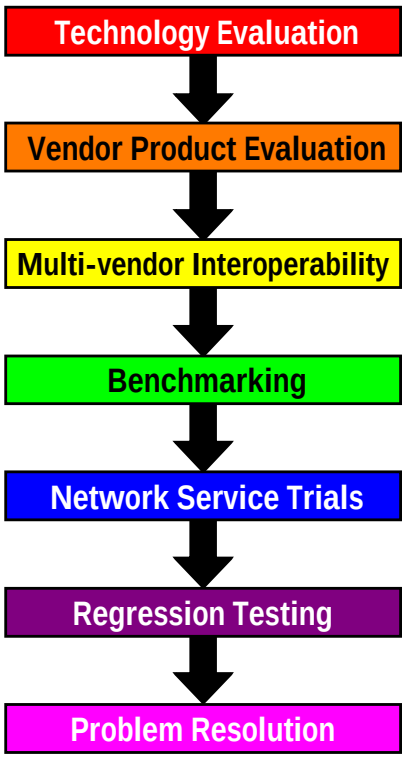

Figure 1: Seven stages of network technology adoption

Stage 1—Technology Evaluation

Many Service Providers evaluate new technologies well ahead of adoption to assess the potential business impact and to become familiar with future challenges.

- Will this new technology reduce operating costs or allow me to offer new, revenue-generating services?
- Is it stable and reliable, or is it just "hype"?
- Are available products mature? Do they meet emerging international standards?
- Will these new protocols and devices reliably scale when used in large networks?

Stage 2—Vendor Product Evaluation

Once a business plan is developed, competing products are evaluated for selection. The chosen network equipment must also pass acceptance testing to ensure that it is fit for the purpose.

- Does this device provide the capabilities claimed by the vendor?
- Does the equipment meet my functional and availability needs?
- Will it be robust when overloaded?
- Will it scale as my network grows?
- Will it be able to meet customer needs or SLAs?

Stage 3—Multi-vendor Interoperability

Interoperability is an important part of product evaluation and deserves to be mentioned separately.

- Do my vendors use incomplete or "enhanced" (proprietary) versions of protocols? Have they interpreted the protocol standards differently?
- Will a new device work with my existing legacy network equipment in my own unique network configuration?
- Will it work with my interconnect partners' equipment?
- Can different vendors' protocol and traffic management configuration settings cause unexpected behavior or network instability?

Stage 4—Benchmarking

Another important aspect of equipment selection that deserves special attention is benchmarking, which addresses performance requirements.

- How do I compare performance of equipment from different vendors under realistic conditions, using my own network configuration?
- How do I simulate large networks with hundreds or thousands of peers? How do I simultaneously emulate multiple routing protocols, set up MPLS tunnels, and generate realistic traffic streams on each tunnel?

Stage 5—Experimental Networks and Service Trials

Following equipment selection and acceptance, an experimental network is often built to test services in a controlled environment that closely mimics the "live" network and customer network traffic. Similarly, a new service may be deployed first in a "trial" within a limited region or with a limited set of customers.

- Before deployment, I want to trial my new equipment. How do I simulate my network and test new services?
- Will my network scale to cope with the demands of new multicast, VLAN, and VPN services?
- If I implement a new traffic management feature, will my network perform as expected? What will be the impact on OSS systems?
- Can my network cope with malicious "Denial-of-Service" attacks from hackers—including attacks on one customer or on the whole network?

Stage 6— Network Upgrades and Regression Testing

After initial deployment, networks are frequently upgraded to expand capacity, to implement new cost-saving features, to be able to create new services, and to eliminate bugs. Unfortunately, each new software or hardware upgrade introduces new risks.

- Will the new switch/router software introduce bugs that are unique to my multi-vendor network configuration?
- Will my existing services still work?
- Are OSS systems providing correct information?
- Has performance been increased or reduced?

Stage 7— Problem Resolution

When a network problem is discovered that cannot be resolved in the field, it is brought back to the lab to test in isolation or within an experimental network. This process is sometimes called "Tier 3 Troubleshooting" in North America.

- When I discover a real network problem, how do I solve it without bringing down my whole network?
- Can I diagnose the *cause* of the failure by simulating it in a repeatable manner in my lab?
- How do I determine the cause of performance bottlenecks that threaten the ability of my network to meet SLAs?
- How do I show my vendor the problems and limitations that I find?

Testing before deployment

The old adage that "prevention is better than cure" definitely applies to network outages. Let's take a look at how testing can help prevent future problems. For each stage of network technology adoption, we will present the following:

- Real-life case study: an actual outage or problem event that occurred
- Relevant test scenario: a test scenario that may mitigate future problems
- Test equipment considerations: important or novel test equipment features that will shorten your test cycles and increase your confidence in the availability of your network services

Stage 1—Technology Evaluation

Case study

One of the largest U.S. carriers' early adoption of "broadband convergence" and Voice-over-IP for gaining a first-mover advantage did not live up to everyone's expectations. The carrier endeavored to converge ATM, Frame Relay, and voice onto an IP backbone, and to offer voice and Internet access over a DSL or T1 connection for both residential and business usage. Although the carrier used off-the-shelf Cisco and Nortel equipment, it was reported that a great deal of customization would have been required to make everything work—according to some, voice services never had the reliability. Early network equipment had problems and was not yet ready for unifying carrier networks and services. The carrier abandoned plans *after spending $3 billion*. While we cannot presume to judge the success of this or any other specific deployment, the lesson for Service Providers is clear: the capabilities and promises of new technology must be tested. (For more information, see reference [3] on page 10.)

Test scenario

Rapid and automated provisioning of bandwidth for network services is currently a hot topic for Service Providers. The emerging OIF UNI and IETF GMPLS recommendations promise to make this dream a reality, offering reduced operating costs and faster provisioning of customer services. Although standards are still evolving, some Service Providers have taken the plunge and deployed early Intelligent Optical Network equipment.

To test whether Intelligent Optical Network equipment is mature for your application and for interworking with your legacy network equipment, follow these general steps:

1 Connect test equipment to the Intelligent Optical Switch or Edge Router.

2 Verify its claimed capabilities by emulating the UNI protocol.

3 Run conformance test suites to verify conformance to early OIF UNI or GMPLS protocol specifications (see Figure 2 below).

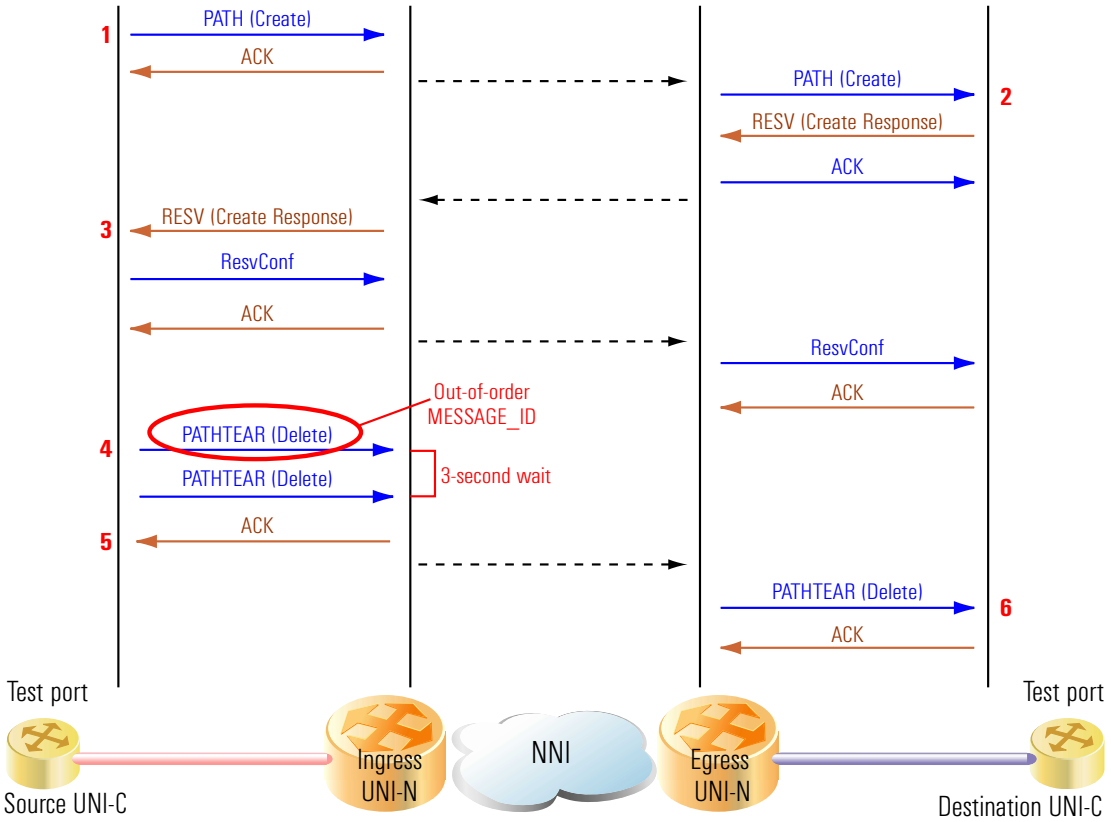

Figure 2: Example of conformance testing

Test equipment

The tester should offer the following features:

- Broad coverage of existing and emerging protocols, including full protocol emulation and a large range of conformance test suites to comprehensively test all aspects of leading-edge network equipment (e.g., MPLS/GMPLS signaling, routing protocols, and security and tunneling protocols, such as L2TP and IPSec).
- Broad range of interfaces, including "legacy" ATM interfaces and high speed 10Gb/s POS and Ethernet interfaces so that you can connect to any port of the system under test.

Stage 2—Vendor Product Evaluation

Case study

On October 28, 1998, AT&T lost 400 T3 lines. Following the initial fiber cut outage, a second outage was caused by routing update traffic. It appeared that every Lucent router nationwide was too busy relearning network routes to actually forward traffic! The resulting routing update traffic effectively "pegged all routers at 100%", bringing about the second outage. Traffic had to be isolated from the backbone. (For more information, see reference [4] on page 10.)

Test scenario—Route convergence speed

Some routers are not robust under high route-update loads, which can lead to major outages during network transients. To measure route convergence speed and verify stability under large routing table updates, follow these general steps:

1 Simulate a large network and take out a primary link. (The router under test must quickly learn the new route—some routers take a long time, or even stall.)

2 Measure route convergence time—i.e., the time between the initial update message and full throughput on the secondary port (see Figure 3 below). Note that some testers only estimate route convergence time from packet loss.

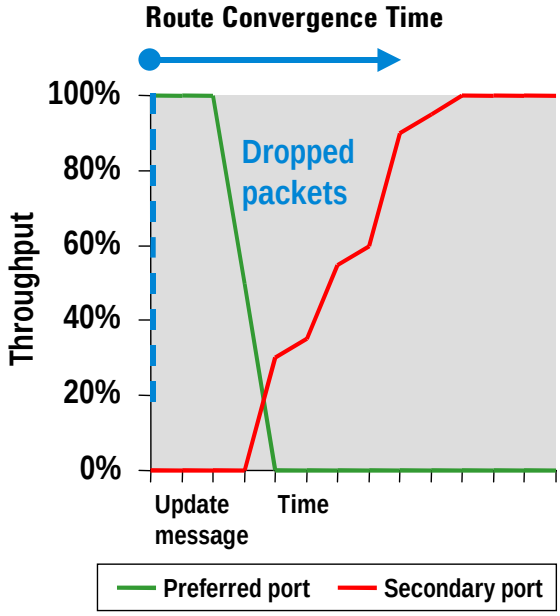

Figure 3: Measuring route convergence speed

Test equipment

The tester must offer tight integration of traffic forwarding, MPLS, and routing protocol emulation. This makes realistic network simulation fast and easy. Some testers require the user to manually send multiple routing and LSP requests and set up many tunnels with traffic, which can waste valuable time.

Stage 3—Multi-vendor Interoperability

Case study

On August 7, 1996, AOL's Internet service went down for 19 hours. Six million customers were locked out. The cause was a glitch in new router software. Vendors claimed it was a "multi-vendor issue", and $2 million was credited to customers. (For more information, see reference [5] on page 10.)

Test scenario—Verifying interoperability and standards compliance

Each device should be tested separately using commercially available conformance test suites. It is important to check for vendor differences in protocol interpretation, which can cause network instability, and to determine whether vendors are using incomplete, out-of-date, proprietary, or "enhanced" versions of protocols. The general test steps are as follows (see Figure 4):

1 Interconnect routers from two different vendors.
2 Inject simulated traffic and routes from test equipment.
3 Monitor and debug protocol interchange between the two devices.

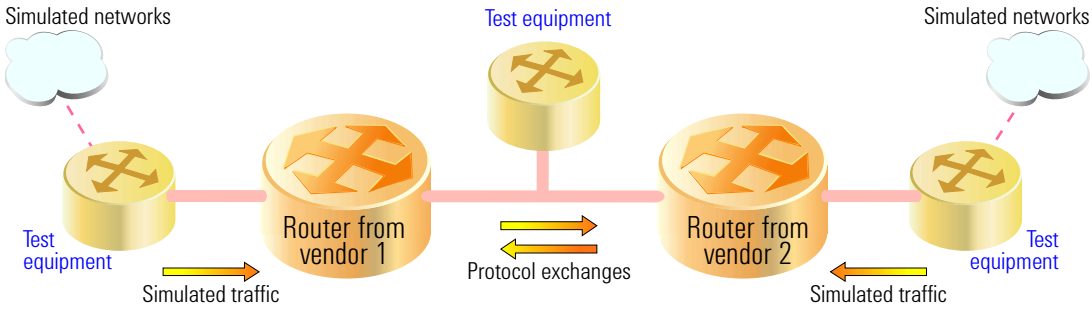

Figure 4: Verifying interoperability

Test equipment

To successfully test multi-vendor interoperability, test equipment must have a large range of conformance test suites and protocol decodes for analysis of interoperability problems. This allows you to comprehensively verify compatibility with new and legacy systems.

Stage 4—Benchmarking

Case study

In May, 2001, a router data flood in Fasthosts' network provider, British Telecom, caused Fasthosts' customers a 36-hour service outage. A routing software bug was thought to be a possible cause, and Fasthosts was blamed for having inadequate redundancy. The router that announced Fasthosts routes using BGP (Border Gateway Protocol) had to be modified to stop the data flood. (For more information, see references [7], [8], and [9] on page 10.)

Test scenario—Measuring traffic capacity

To measure traffic capacity, follow these general steps:

1 Generate up to 2,000 traffic streams per port between thousands of test ports to stress the router forwarding engine and switch fabric resources (see Figure 5).

2 Configure initial traffic bandwidth, then decrease bandwidth to under-subscribe the connection.

3 Finally, increase traffic to over-subscribe the connection.

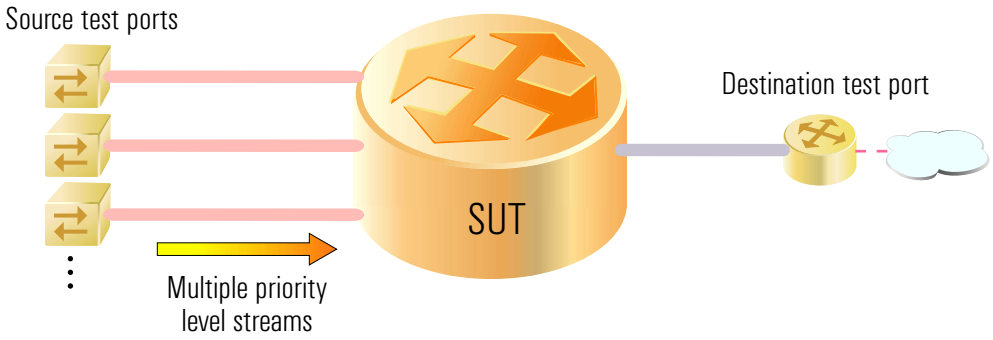

Figure 5: Measuring traffic capacity

Test equipment

Test equipment must be able to control generated traffic bandwidth interactively without interrupting the test. Many testers introduce generated jitter (inter-arrival time bursts) when traffic levels are changed (see Figure 6).

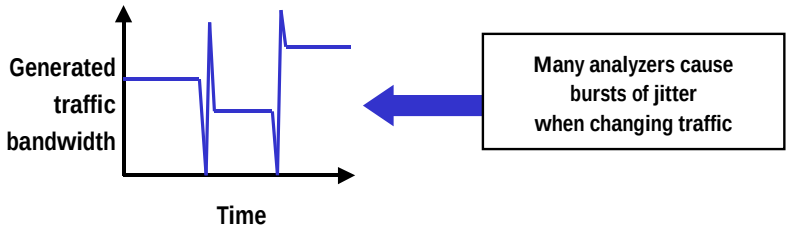

Figure 6: Jitter induced by test equipment

Stage 5—Experimental Networks and Service Trials

Case study

On December 8, 2001, Telstra's data network slowed to a crawl for an hour due to two Denial-of-Service (DoS) attacks. These are malicious attacks by hackers that can severely disrupt services, either by attacking another network host or by attacking the whole network. DoS attacks regularly cause outages and service downtime. (For more information, see reference [10] on page 10.)

Test scenario—Simulating DoS attacks

You can verify the ability of network equipment to resist malicious DoS attacks by testing in a safe, "out-of-service" test environment (within an experimental network or evaluation lab). Figure 7 illustrates how you can simulate attacks while generating and measuring the performance of "background" user traffic. Ideally, user traffic performance should not be impacted by the attacks. You can mitigate future network outages and service downtime by reconfiguring network devices to cope with DoS attacks.

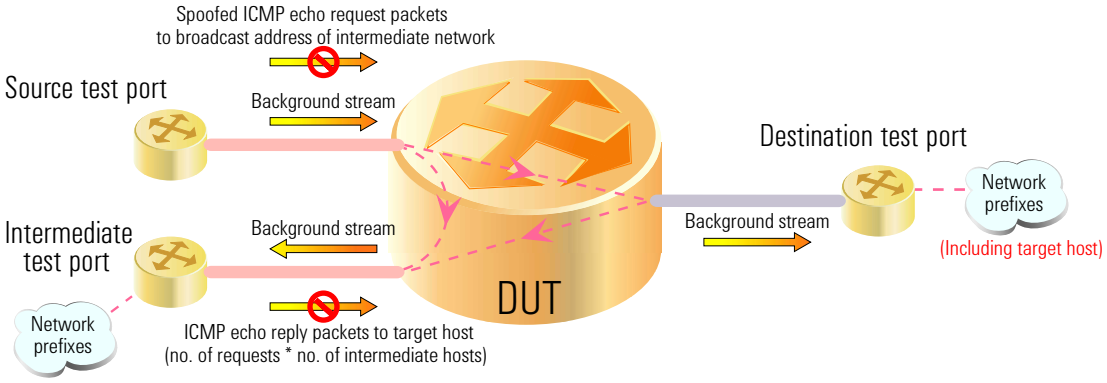

Figure 7: Example of simulated DoS attack

Test equipment

Test equipment must provide automated software that can simulate a large range of Denial-of-Service attacks.

Stage 6—Network Upgrades and Regression Testing

Network upgrades are the most common cause of outages. According to Jeff Moore, Senior Telecom Analyst at Current Analysis, Inc., "Most network switch problems occur during software changes." (See reference [11] on page 10.)

Case study

On April 13, 1998, AT&T's Frame Relay network went down for 26 hours. The cause was faulty software during a switch software upgrade—a problem unique to the specific configuration of one switch! (For more information, see reference [12] on page 10.)

Test scenario—Regression testing

Developing a regression test plan to cover the range of functional capabilities, performance requirements, and services offered by the network equipment is essential. Common test cases should be automated so they can be run quickly and accurately during every major upgrade of network software or hardware (see Figure 8 below).

Test equipment

The test vendor should provide a journal of relevant test cases that you can use as building blocks to construct your test plan. Vendors should also provide a scripting environment that allows customization and development of applications to automate your test plan quickly and easily. Conformance test suites can provide a fast and automatic check of new network software.

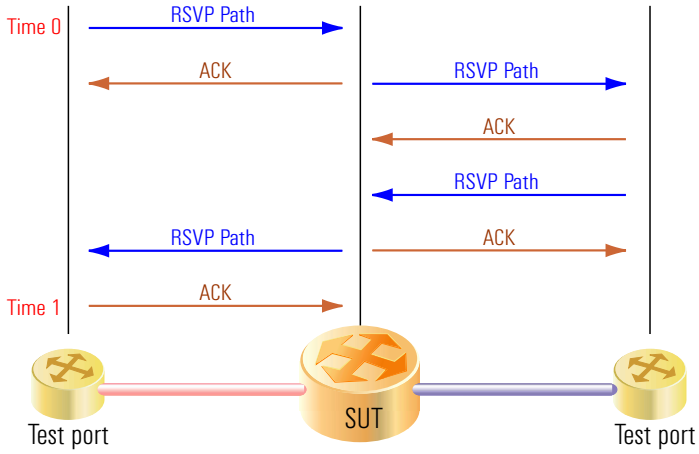

Figure 8: Verifying that a router software upgrade will not introduce network bugs

Stage 7—Problem Resolution

Case study

On November 21, 2000, an Australia-Singapore cable cut in Telstra's network caused routing loops on core routers. Outages lasted over 10 hours before the problem was found and fixed. (For more information, see reference [13] on page 10.)

Test scenario—Diagnosing performance bottlenecks

In the lab you can simulate a performance issue found within a live network, as follows:

1 Connect a tester to an IP router and simulate user traffic, MPLS signaling, and routing protocols on multiple streams across multiple ports.

2 With the tester, monitor QoS on each stream simultaneously in real time to check for performance degradations.

3 When a performance bottleneck is detected, you can then analyze both the logged performance measurements and the captured data file.

Test equipment

The tester must offer performance-based capture triggers to enable data capture to be triggered when a performance metric exceeds a user-settable boundary— for example, when latency exceeds 100 ms. The cause of performance bottlenecks can often be deduced by examining packets captured on either side of a performance threshold crossing (see Figure 9).

The ability to emulate three or more protocols simultaneously is valuable for solving complex problems involving services such as VPN or Multicast. Simulation of the problem in the lab may require the tester to emulate MPLS signaling, VPN service protocols, IGMP / PIM-SM multicast protocols, and several routing protocols at the same time.

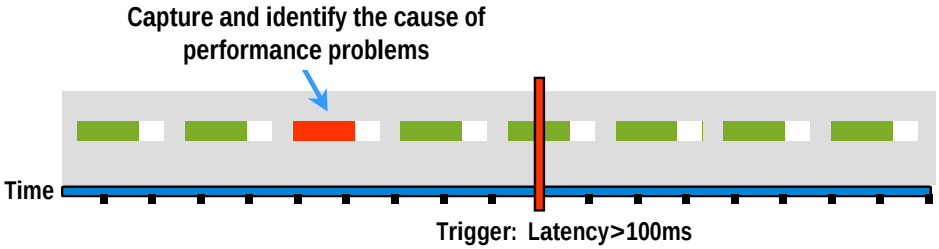

Figure 9: Example of a performance-based capture trigger

Summary

There have been many documented examples of lengthy network outages and costly service downtime. If you are interested, follow the links in “Further Reading” below to read about other outages.

While vendor testing of new network equipment is often extensive, it is not a replacement for Service Provider testing during the seven stages of network technology adoption—technology evaluation, vendor product evaluation, interoperability, benchmarking, experimental networks and service trials, network upgrades, and problem resolution.

Prevention is better than cure. That is, it is more cost effective to find network problems prior to deployment than it is to suffer network outages. Some outages are inevitable, but Service Providers can mitigate many outages by testing new network equipment and services before deployment.

The single most common cause of outages stems from network upgrades, typically from new releases of network equipment software. A well-designed regression test plan should retest both functionality and performance. To simplify equipment upgrades and reduce the time to market of new services, many Service Providers automate their regression testing, using both commercially available test suites and customized software applications.

References

[1] The Risks Digest, Volume 19, Issue 72, May 1998 – http://catless.ncl.ac.uk/Risks/19.72.html#subj6.1

[2] Sciodata – http://www.sciodata.com/home.asp

[3] Light Reading, December 27, 2001 “Whatever Happened to Sprint's ION?” – http://www.lightreading.com/document.asp?doc_id=10558

[4] The Risks Digest, Volume 20, Issue 5 – http://catless.ncl.ac.uk/Risks/20.05.html#subj4

[5] CNet News.com Tech Site – http://news.com.com/2100-1023-220635.html

[6] Agilent Technologies – The Journal of Internet Test Methodologies

[7] The Register, June 27, 2001 “Fasthosts outage outrage” – http://www.theregister.co.uk/content/23/19334.html

[8] The Register, June 5, 2001, “Fasthosts redundancy redundant” – http://www.stoporbs.org/content/23/19420.html

[9] Net4Nowt public comments page; Net4Nowt is a news service and directory of UK Free ISPs – http://www.net4nowt.com/comments/991256717,32280,.shtml

[10] Matrix News, Issue 11.3 – http://www.net4nowt.com/comments/991256717,32280,.shtml

[11] ITworld.com, “Network Outage Hits AT&T's ATM Users”, Computerworld, February 26, 2001 – http://www.itworld.com/Net/2318/CWSTO58077

[12] AT&T Press Release, April 22, 1998, “AT&T announces cause of frame-relay network outage” – (http://www.att.com/press/0498/980422.bsb.html

[13] Linux South Australia User Group Mailing List Archive, November 2000 – http://www.linuxsa.org.au/mailing-list/2000-11/834.html

Further Reading

- For more examples of lengthy network outages that have been made public, see http://www.cctec.com/maillists/nanog/historical/0101/msg00764.html or http://www.matrix.net/publications/mn/mn1103_microsoft.html
- According to a Communications Week user survey, copper-based networks average 2.3 network outages per month, at an average cost of over $19,000 – http://www.fols.org/pubs/fiber_to_the_desk.html
- A 1998 Merit Networks study of IP networks that included three major ISPs found that 10% of routes had less than 95% availability, and that only 30% of outages were repaired within one hour. “Experimental Study of Internet Stability and Wide Area Backbone Failures,” by Craig Labovitz and Abha Ahuja, Merit Networks, 1998 – http://citeseer.nj.nec.com/labovitz98experimental.html and http://www-unix.ecs.umass.edu/~lgao/class/routing/lect11.ppt
- A second outage to Telstra's network that was extended by a DoS attack is described in http://www.matrix.net/publications/mn/mn1103_microsoft.html

March 15, 2002
5988-6087EN

Agilent Technologies