

Understanding How Applications Perform on the Network— Even Under Adverse Network Conditions

Enterprises are under significant pressure to make sure they get the best possible performance from network-connected services and applications. Whether an enterprise is a manufacturer developing software for connected devices or a corporation adopting business solutions, the organization's engineering and IT executives need assurance that the forthcoming applications will always work as expected. The implications of a compromised product are too big to ignore: Depending on the application and use case, the risks include productivity losses, development cost increases, dissatisfied customers, lost revenues, even adverse associations to a company's brand.

Ensuring product performance is complicated and challenging, however, because organizations must engineer their applications with full understanding of network conditions. Yet most applications traverse a variety of networks that are beyond the organization's control. Undertaking a trial deployment to evaluate the impacts is costly and inherently limited because a trial is unlikely to encounter the full range of impairments that can occur on a network.

What's Inside...

- ▶ Introduction (pg. 1)
- ▶ The Service Context (pg. 2)
- ▶ Networks are Not Perfect (pg. 3)
- ▶ The Network Emulation Approach (pg. 5)
- ▶ Mitigate Performance Issues (pg. 6)
- ▶ Make Sure Your Next Application Rollout is a Success (pg. 7)
- ▶ About Interworking Labs (pg. 8)

Fortunately, enterprises today can emulate real-world networks in their labs to understand how their products perform under routine and adverse network conditions. This paper summarizes the types of adverse conditions that can occur and recommends that organizations incorporate network emulation tools into the testbed to reveal how their current and planned applications behave in the presence of network problems.

Organizations can use findings from these tests to define the service levels their applications need or guide negotiations for service-level agreements (SLAs). This approach will also ensure services perform as needed under a broad range of conditions. Businesses that use the approach can shorten time-to-market with high-quality applications and deploy applications with confidence while reducing costs and risks.

The Service Context:

Applications are Very Demanding and Sensitive to Network Problems

The Internet is becoming more complex and the services and applications that run on the network are more demanding than ever before.

A decade ago, networks carried substantially less traffic than they do today, applications required less bandwidth, and capacity limitations were not much of a concern. Also, applications were fairly basic. Customers used the network for Internet browsing and non-conversational (one-way) streaming audio and video. The applications could deliver a good user experience even in the context of impairments that can occur on the Internet, such as packet transmission delays.

Today the network environment is much more demanding. In recent years networks have advanced continually to accommodate increasing traffic demands and the expectations are not letting up. By 2019, global networks will need capacity to support 142 million people streaming high-definition video over the Internet simultaneously all day, every day, according to Cisco. Every two minutes, the gigabyte equivalent of all movies ever made will cross the Internet. Broadband speeds are also increasing and will likely double by 2019 compared to speeds offered in 2015.¹

As network performance and demands increase, usage patterns evolve accordingly. Today, for example, companies are deploying applications that place new and significant demands on underlying networks. The applications

are also extremely sensitive to nuances in network performance. For example, storage-area networks (SANs) that keep corporate data near a business for time-sensitive applications, conversational tools such as voice-over-IP (VoIP) services, and real-time applications such as multiplayer online games are directly influenced by packet transit time and other impairments that can affect packets during transmission.

Networks are Not Perfect: Transmission Problems Will Occur and Impact Application Performance

We expect a lot of our networks but we shouldn't take them for granted. While the technology industry and standards organizations have engineered networks to meet current and future communications needs, there is no such thing as a perfect network. Ironically, the laws of physics and mathematics that establish the speed of electrical pulses on wired connections or the speed of light on fiber optic cables can also impose random limitations, such as noise on cables or hardware or software failures in routers and switches.

There are additional sources of imperfection that affect services and the user experience. Packets can be lost or delayed by transient congestion in the Internet's switching elements. Packets can be lost, replicated, or transmitted out of order as a result of routing changes. Packet reordering can occur during load balancing—when traffic between a pair of routers is sent over parallel links to avoid congestion or a network failures. Packets can become corrupted as well. Depending on the type of impairment that occurs, end users could experience service interruptions, diminished throughput, or poor video or voice quality, among other performance problems.

These conditions tend to occur in bursts that can last a few milliseconds or a few minutes, however longer periods are not atypical. The impairments do not occur very often in the core of the Internet, which is expertly engineered, always state-of-the-art, and closely monitored 24/7, however core network problems do occur occasionally. Most impairments occur in the periphery of the Internet, where most packets enter or leave the network. Edge-of-network impairments are usually encountered when traffic passes through overloaded Internet exchange points or older, under-provisioned or poorly monitored devices such as routers that connect an internal local area network (LAN) with an external wide area network (WAN) or the Internet.

Can an organization's own network suffer these impairments? The answer is yes. Any network that has a few switches and routers, and especially any network that includes off-campus connections, is likely to have service impairments. In fact, smaller networks are more likely to have impairments than larger networks because larger networks are typically monitored 24/7 by network operations centers (NOCs) that constantly observe network conditions to spot problems in real time, pinpoint the causes and initiate repairs.

Impairment	Definition
Packet Loss	The disappearance of a packet that should have been transmitted.
Packet Delay	The time it takes for a packet to travel from its source to its destination.
Packet Delay Variation (Jitter)	The variation in packet delay, when some packets reach the destination faster than others. Can be accompanied by packet loss.
Packet Duplication	When a packet is duplicated and multiple identical copies are transmitted.
Packet Reordering	When a network delays some packets, but not others, so they arrive out of sequence at the destination.
Packet Corruption	The contents of a packet are damaged but the packet continues to flow to the destination.

What are the effects of packet impairments?

- ▶ Applications hanging or failing
- ▶ Diminished throughput
- ▶ Less responsive web browsing
- ▶ Poor voice quality
- ▶ Poor video quality
- ▶ Pixelated video images
- ▶ Video images and sound out of sync
- ▶ Intermittent service interruptions
- ▶ And many more...

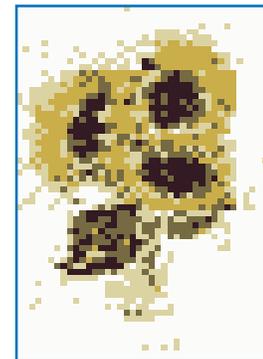
The Network Emulation Approach:

How to Identify the Impacts of Adverse Network Conditions on Application Performance

Assuming every network will have some type of impairments, enterprises need a practical method they can use, with confidence, to identify how their applications will perform under routine and adverse conditions.

The recommended approach is to build a network emulation test bed to emulate a range of network conditions that the organization's customers might encounter in real-world scenarios. The enterprise can use the test bed to introduce adverse network conditions, evaluate how the applications behave in the presence of each condition, determine the range of impairments that the company's applications can tolerate, and define the level of service the application requires for reliable, effective use in the market.

Engineers and IT executives often assume that testing is easy and straightforward, but a variety of implicit challenges must be considered for a successful implementation that yields meaningful results. For example, some organizations will build a network in the lab to replicate a couple nodes in the data center, but tests performed on this type of system will not reveal problems that can be introduced when high



numbers of smart phone users access the network. Other organizations will test their applications on a “perfect” network they build in the lab, but such networks typically do not represent reality because the lab's nodes are installed and configured correctly, in close proximity. Such approaches can be expensive, inconvenient and costly to adapt if alternative configurations must be considered. And unfortunately, neither approach will accurately reflect network conditions found in the real-world or the behavior of a full-scale network.

Another option is to use mathematical simulations and models. These methods take a great deal of expertise to design, implement and evaluate, however, and can be misleading: It's very unlikely that software deployed on devices in the market, for example, will perform with the mathematical precision of simulated conditions.

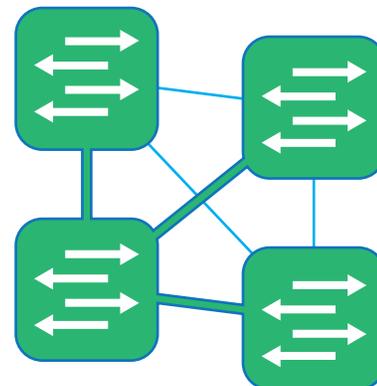
The best approach is to employ a network emulator that can conveniently produce a variety of network impairments in an application test bed, under controlled and

repeatable conditions, so companies can evaluate the performance of proposed applications under real-world conditions. The tools should have not only the ability to induce the types of impairments described in this paper; the tools should also have the capability to induce and expose implementation problems that open up opportunities for network attackers to break in or create denial of service attacks.

Mitigate Performance Issues: Use Test Bed Findings to Define Service Needs and Guide Service Level Agreements

Organizations can use the findings from their network emulator test bed evaluations to determine the conditions their applications require and to establish service-level agreements (SLAs) with their network providers to ensure their networks perform as needed.

The sensitivity of applications to network impairments can vary widely depending on the types of applications deployed and their performance requirements. Certainly, organizations can and should optimize their applications to perform as well as possible under a range of conditions, but high-performing applications must often operate on a network that is engineered to guarantee high-quality packet delivery. Real-time applications such as VoIP, for example, will be more susceptible to impairments than email and therefore many enterprises will want to come up with appropriate service-level definition and SLAs for their VoIP services.



Spanning Tree Protocol

A pragmatic approach for pinpointing these needs is to inventory the applications running on the network, identify each application's performance requirements, evaluate how the application performs under adverse network conditions, and define the service level needed to appropriately support the applications. The service level definitions should establish minimal performance requirements for the key variables that will impact service quality, such as packet delay and jitter. The organization can also specify performance values for packet loss, duplication, reordering and corruption.

Enterprises can use service-level definitions to guide their network deployment decisions and make the best use of their investments in network infrastructure. The definitions will also enable companies to avoid unnecessary time and money over-engineering their networks to provide service levels their applications don't need.

With the service definitions established, organizations can also negotiate SLAs from service providers guaranteeing the network will provide the specified level of service. Impairments that do not impact services may be safe to ignore. It is important, however, to review SLAs as new devices or applications are added, old ones are removed, and as overall traffic demands and usage patterns change. The introduction of new infrastructure, such as software defined networks, can also change network conditions and operational characteristics that impact application performance and SLAs.

Companies can also use findings from their network emulation testing to confirm or refute claims from network equipment vendors that certain classes of applications require new equipment deployments. Often these claims are unwarranted or based on a vendor's misunderstanding of the application's actual requirements, and network emulation testing and analysis can help clarify needs to avoid investments in unnecessary equipment.

Make Sure Your Next Application Rollout is a Success:

Learn More about Pre-deployment Testing

Enterprises are under considerable pressure to ensure that the applications they develop or deploy will perform as promised and withstand the vagaries of the Internet. Fortunately, network emulators, integrated in an application test bed, provide a convenient and cost-effective method that organizations can use to pretest their applications under a wide range of network conditions to define needed service-levels and establish SLAs. The approach enables companies to introduce applications with confidence while reducing costs and risks.

To learn more about the uses and benefits of test beds for pre-deployment evaluations, contact InterWorking Labs.

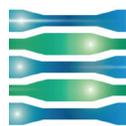
About InterWorking Labs

InterWorking Labs (IWL) helps developers and engineers understand how well a device or application will perform under adverse network conditions, not just perfect conditions. The developer can then improve the embedded system or network application to optimize its performance. Performance can range from user-perceived response time to robustness in the face of pathological packets. Maxwell® and KMAX™ network emulators offered by InterWorking Labs enable software developers, network managers and test engineers to create test beds in their labs to guide their deployment decisions. For more information about Maxwell and other IWL solutions, visit <http://iwl.com/>.

¹ Cisco. "The Zettabyte Era—Trends and Analysis," May 2015.

http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/VNI_Hyperconnectivity_WP.html

Want to learn more?

The logo for InterWorking Labs, featuring a stylized graphic of three horizontal, wavy lines in shades of blue and green.

Kings Village Center #66190
Scotts Valley, CA 95067
iwl.com
+1.831.460.7010
info@iwl.com