Predicting and Managing Network Impairments

Table of Contents

Industry Drivers

As networks evolved to focus on cost and scale advantages, Ethernet technologies came to the forefront. Like the time-division multiplexing (TDM) networks of old, today's Ethernet packet transport networks (PTNs) have deficiencies that negatively impact service quality and overall user experience, and also bring a whole host of new reliability and latency problems.

Network imperfections, or impairments, can put network performance at risk and thus jeopardize key aspects of a company's operation. Subscriber volatility is fueled by packet delay, jitter, packet loss, and other problems that impair and disrupt vital services and affect subscribers' quality of experience. These network imperfections are unavoidable impediments that network designers and application authors need to mitigate.

Network imperfections are unavoidable impediments that network designers and application authors need to mitigate.

Figure 1. Some applications and services are hyper-sensitive to network impairments that cause service disruptions

Network equipment manufacturers (NEMs) must ensure their products do not create unnecessary impairments and can withstand and even compensate for a certain level of impairment that is naturally introduced into the network infrastructure. They need to validate their devices for robustness in the presence of impairments to compete with other vendors and to ensure customer satisfaction.

Service providers, enterprises, and government agencies that own or run networks work hard to prevent the service quality problems caused by unexpected network impairments. To keep their customers happy and to maximize their capital investments, these network operators must ensure that their core network, consisting of many traditional and application-aware devices, maintains proper QoS for all voice, video, and data traffic.

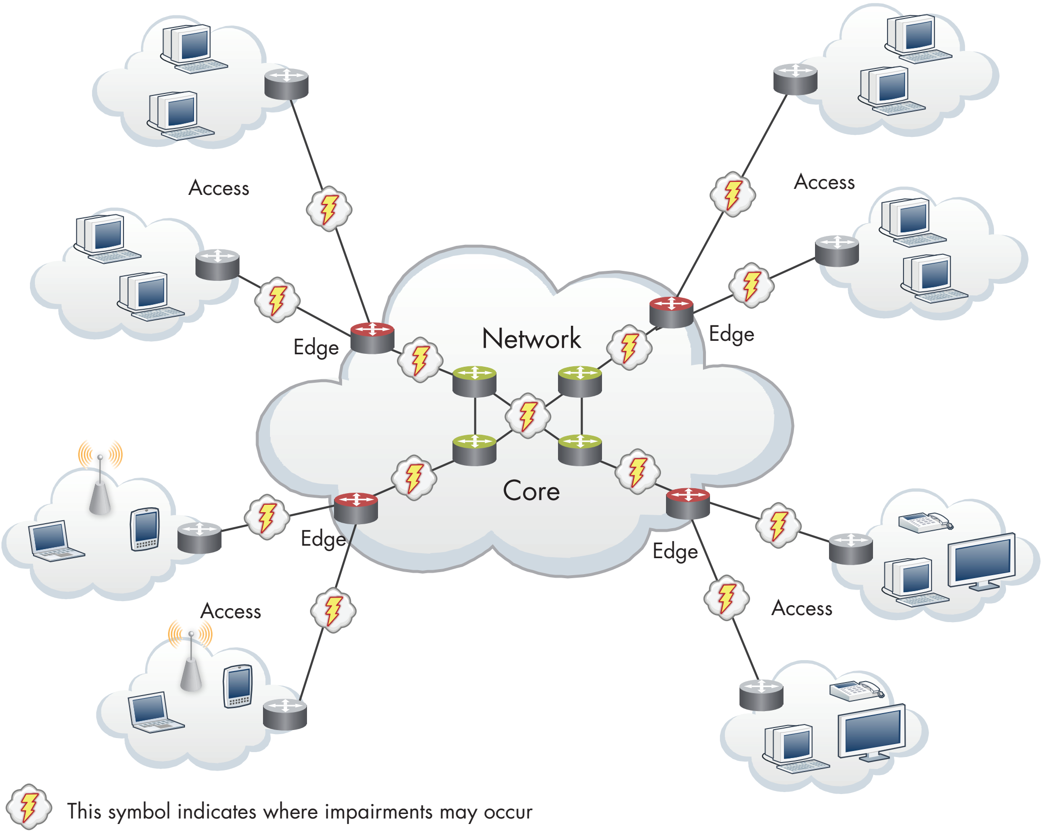

Figure 2. Impairments occur in any part of your network and can be far-reaching

Since real-world networks do not behave in a deterministic manner, impairments that play havoc with both control and data packets passing through networks must be discovered and rectified, especially for quality-sensitive service infrastructure such as:

- Financial applications
- Voice and video traffic
- Cloud and network applications
- Enterprise applications
- WAN optimization
- Carrier Ethernet
- Routing and MPLS
- Satellite networks
- SLA planning and validation
- Data centers
- Network security

Impairment Problems

As we've all experienced, working networks do not behave in a deterministic way. Caused by cumulative random events, impairments affect traffic packets that are traversing a network differently at any given moment. Most of us have experienced slower Internet upload/download speeds in the afternoon when the kids in our neighborhood get home from school. Distance of the path dynamically built for each packet as it traverses a cross-country network can impact application performance. And what about the media frenzy that happens when there is a wide-spread service disruption? Impairments are fickle, coming and going with whatever is trending.

It's not the intent of this paper to detail each type of impairment, but we'll briefly discuss the top three that have the most impact on quality-sensitive services.

Packet Delay

Packet delay, also known as "latency", is a measurement of how much time it takes for a data packet to get from one point to another. Although a measurement of zero delay is not seen in production networks because it takes some amount of time for the packet to travel from its source to its destination, a low packet delay number is desired for optimum network and application performance. Propagation, router/switch processing, and storage delays are normal contributors to the expected delay. However, adverse network conditions like queuing delays on the intermediate network elements also contribute to packet delays.

Packet Loss

Packet loss is a measure of the number of packets sent over a network that fail to reach their destination. This can result in noticeable performance issues. In addition, a decrease in the throughput is caused by some transport protocols such as TCP, which have a mechanism to ensure reliable delivery of packets, requiring the retransmission of missing packets. Causes of packet loss may include multi-path fading, channel congestion, in-transit rejection of corrupted packets, faulty hardware, drivers, or routing routines.

Delay Variation

Delay variation, sometimes referred to as jitter, is the measure of variability of delay values over a period of time. The delay experienced by traffic is not a static value but usually varies due to random events such as fluctuating loads on the network infrastructure. For example, in the morning when a lot of people log on to the network and start accessing network resources versus the evening hours when people retire to bed and network usage goes down. This dynamic nature of network traffic affects the delay that packets experience as they traverse the network. Packets suffering from delay variance will end up arriving at the destination out of sequence or may even be dropped by the receiving devices. This has a negative impact on the performance of various voice and video applications.

Since all networks contain impairments, the relevant question is, "how much of an impairment is too much for my network, service, or device?"

Right-Sizing Impairment

Since all networks contain impairments, the relevant question is, "how much of an impairment is too much for my network, service, or device?"

The impact of bandwidth, latency, jitter, and other factors is referred to as the quality of experience (QoE) and is a reflection of how satisfied end-users are with the services they receive. Each type of service comes with its own requirements. Voice over IP (VoIP) has a very low bandwidth requirement, but requires low latency and jitter. IPTV uses high bandwidth, but is tolerant of moderate jitter. A particular P2P application, on the other hand may have very high bandwidth requirements and can sustain high latency and jitter.

Network operators must ensure that the QoE of their services “feels right.” VoIP calls must sound as good as land-line service, IPTV must not have pixelated, blurred, or frozen frames, and high-speed Internet services must appear responsive. Subscribers have a very low tolerance for these types of defects; they’ll quickly switch to competitors. In addition, business customers often demand specific performance in their service level agreements (SLAs). If those SLAs are not achieved, then the network operator has to pay penalties.

Impairments are always present within devices, systems, and networks to a certain degree. To ensure interoperability between devices and minimize service degradation across long distances, it is important that there are limits set on the maximum level of the most relevant impairments present at an output interface and the minimum level that can be tolerated at an input.

Testing for these impairments before deploying new devices, services, application, and networks saves time and money.

Standards Bodies

Adherence to pre-determined limits will ensure interworking between different vendor equipment and networks, as well as providing the basis for isolation of problems. These limits have been determined by industry organizations and then implemented in standards for different types of devices, interfaces, and applications. Because there are so many device and interface types, finding the relevant impairment standard for a specific application can be difficult. Here’s a list of some of the standards bodies concerned with the efficient transportation of communications traffic:

- *3rd Generation Partnership Project (3GPP)*
- *Alliance for Telecommunications Industry Solutions (ATIS)*
- *American National Standards Institute (ANSI)*
- *European Telecommunications Standards Institute (ETSI)*
- *Institute of Electrical and Electronics Engineers (IEEE)*
- *International Telecommunication Union (ITU)*
- *Internet Engineering Task Force (IETF)*
- *Next Generation Mobile Networks Alliance (NGMN)*
- *Optical Internetworking Forum (OIF)*
- *Small Cell Forum*
- *TM Forum*

It is imperative for network designers, enterprise architects and application engineers to verify that their solution can withstand the vagaries of network conditions. Testing for these impairments before deploying new devices, services, application, and networks saves time and money that would otherwise be spent in fixing applications and networks post deployment and trying to win customers back.

Impairment Testing Requirements

Pre-deployment impairment testing solutions enable NEMs, service providers, enterprises, and government agencies to ensure resilience to inevitable network impairments. They offer a controllable, simulated network environment that lets you evaluate how networks and services are impacted under impaired traffic conditions.

The use of test equipment for impairment testing allows for savings in time, effort, and money that would otherwise be spent using expensive network resources to replicate real-life conditions. With an comprehensive impairment testing solution, you can easily emulate a real-life network with a few clicks of the mouse, running network protocols and data over an emulated network that has impairments impacting those packets.

Top requirements for an impairment testing solution should include:

Impairment Test Feature	Function	Benefit
Hardware-based impairment generation	1GE, 10GE, and 40GE line rate impairment for all frame sizes with no packet loss, built-in traffic classifiers for primary protocols and applications, and ultra-high latency simulation	Provides superior network performance, cost-effective testing, and superior test accuracy with guaranteed line rate performance impairments across all frame sizes.
Seamless integration with traffic generation, protocol emulation, and analysis from a single user interface	Impairment built into L2/3 and L4-7 test applications, offering complete protocol and data plane integration with impairment	Faster time to test, lower costs, and ease of use. No need to learn new GUI and switch back and forth. No need for extra pizza boxes.
Industry-leading performance and port density	High-density 1GE, 10GE, and 40GE support	Maximize ROI with a system that supports your current and future network infrastructure
High levels of latency	Can emulate long network distances and corresponding network latency at line rates. Look for 600ms delay with 10GE port pair, 6s with 1GE port pair, and 500ms delay with 40GE port pair.	A precise ability to create and measure time- and distance-based performance
Realistic, high-scale WAN emulation	Emulate large numbers of network clouds with various link speeds	Enables precision testing of expansive networks
Flexible classifiers that allow for multiple custom selectors to select traffic to be impaired	Offers greater granularity in the targeting of particular subset of packets to be impaired	More precise impairments that emulate live networks by varying the way traffic is impaired

Look for seamless integration with traffic generation, protocol emulation, and analysis from a single user interface.

To test and measure impairments, traffic ports are connected to the devices under test (DUT), which are then connected to impairment ports that emulate the WAN. The applications that control both the traffic ports and the impairment ports should be accessed through a single graphical user interface (GUI).

Impairment profiles are set that determine the treatment of certain packets. The most common impairment values should be available as pre-set menu items, while also making it simple to enter custom configuration values. Specific functions to look for in a robust impairment solution:

Rate limit: Rate limiting allows the user to emulate the behavior of intermediate network devices that may impose rate limits due to either their configured network policies or because they are overloaded and cannot sustain high traffic rates. Rate limit options should include Kb/s, Mb/s, Gb/s.

Delay: Set a delay of n number of microseconds, milliseconds, seconds, or kilometers. The kilometers value automatically calculates the delay, which is useful when the distance between two endpoints in the WAN is a known value. To emulate jitter, the impairment solution should randomly vary the delay applied to each packet using uniform, Gaussian, or exponential distribution. You should also have the ability to set the probably of delay and delay values.

The most important impairment values should be available as pre-set menu items, while also making it simple to enter custom configuration values.

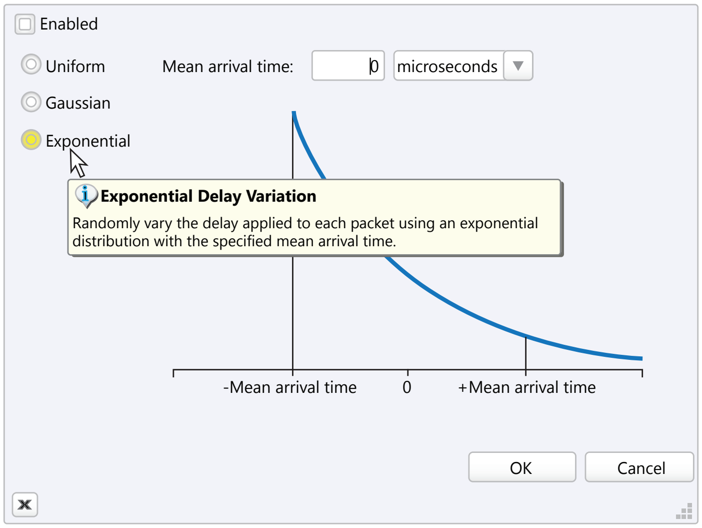

Figure 3. Because jitter impairments derive from randomly-varied traffic delays, test tools must have the ability vary delay distribution

Drop: state the percentage of packets to drop, the number of packets to drop at a time, reorder the packets, and duplicate packets. This allows flexibility in emulating the types of drops you want to see on the network.

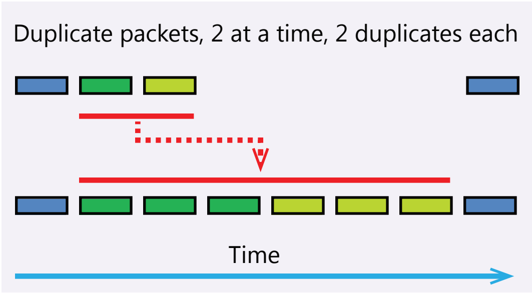

Figure 4. Example of duplicate packets impairment

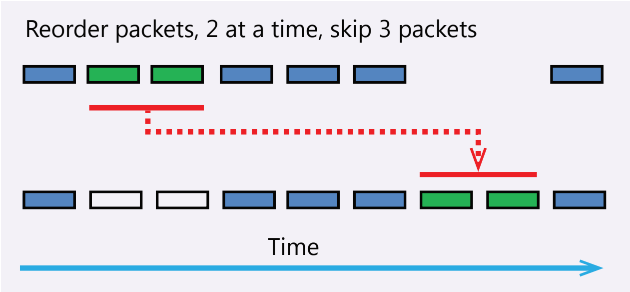

Figure 5. Example of reordering of packets impairment

Links: this controls the directional characteristics of traffic so you can apply impairment in a particular direction from port 1 to port 2, from port 2 to port 1, or both ways. This allows you to configure different impairment characteristics depending on the traffic direction.

Bit Errors: insert bit errors into the traffic stream and even choose to skip octets at the beginning of each packet or the end of each packet. This makes it simple to only insert bit errors into the payload and skip the header and trailer.

Checksums: choose to drop incoming packets with L2 FCS errors or correct L2 FCS errors in outgoing packets.

Be sure your impairment system provides an easy way for you to define the subset of traffic that passes through the impairment ports that will have the set impairments applied to it. Finally, statistics views should show the results of the applied impairments and the effect of the impairments on your traffic.

Be sure your impairment system provides an easy way for you to define the subset of traffic that passes through the impairment ports that will have the set impairments applied to it.

The Ixia Solution

Ixia's comprehensive network impairment test solutions can accurately emulate infrastructure, mobile backhaul, multimedia, security, and storage and application networks, while providing repeatable results that will not fluctuate when the same tests are performed at different times or in different labs. Our network impairment test solutions are comprised of ImpairNet® and Ixia Network Emulators. Ixia also provides the Ixia 3500, an Ethernet time and clock synchronizaiton solution. Please visit our web pages to learn more about how these solutions can help you validate network and application performance:

ImpairNet: 1/10/40Gbe real-world network impairment emulator
http://www.ixiacom.com/products/impairnet

Ixia Network Emulators: Emulate real-world network impairments
http://www.ixiacom.com/products/ixia-network-emulators

Ixia 3500: Validate Ethernet timing/clock synchronization
http://www.ixiacom.com/products/ixia-3500

Conclusion

Network impairments can cause noticeable service disruptions and customer churn. Network operators need to pair acceptable impairment thresholds with each type of service offering and then test prior to deployment to ensure QoE. NEMs must ensure that their equipment does not cause unnecessary impairment and can withstand and even compensate for a certain level of impairment.

Effective impairment test solutions include traffic- and control-plane protocol configuration so you can easily emulate your real-world network and impair the traffic in the same system. These functions go hand-in-hand and should be tightly linked.

Hardware-based impairment solutions, such as ImpairNet, provide robust functions for line-rate impairment for all frame sizes with no packet loss. Testing must include the flexibility to create the wide range of impairments found in real-world networks.

Although out of the scope of this paper, companies should use the results of their impairment testing to determine where to focus efforts to improve performance problems. One view is to keep the network simple; the applications must contain the intelligence to most-efficiently manage traffic flows. Others point out that networks may need to be better provisioned or tuned to handle certain high-bandwidth, time-sensitive traffic. We expect that the answer lies somewhere between the two, with a fine-tuned network and impairment-aware applications.

Ixia Worldwide Headquarters
26601 Agoura Rd.
Calabasas, CA 91302

(Toll Free North America)
1.877.367.4942

(Outside North America)
+1.818.871.1800

(Fax) 818.871.1805
www.ixiacom.com

Other Ixia Contacts
Info: info@ixiacom.com
Investors: ir@ixiacom.com
Public Relations: pr@ixiacom.com
Renewals: renewals@ixiacom.com
Sales: sales@ixiacom.com
Support: support@ixiacom.com
Training: training@ixiacom.com