WHITE PAPER

Plotting the performance landscape for SD-WAN

WHITE PAPER

How Ethernet impairment testing can enable efficient and cost-effective testing of dynamic SD-WAN solutions

Contents

EXECUTIVE SUMMARY

SD-WAN was first introduced to provide a reliable and cost-effective backup to MPLS-based enterprise WAN connectivity. As IP and Internet networks have become more reliable, along with the growth in adoption of cloud services, SD-WAN has evolved and is replacing branch office routers as the preferred WAN connectivity solution.

According to Gartner, 60% of enterprises will have implemented SD-WAN by 2024 compared to just 30% today. At least 30% of enterprise locations will only have Internet WAN connectivity, which is twice the current number. The impact of SD-WAN on the WAN infrastructure market is also clear with Gartner expecting a Compound Annual Growth Rate (CAGR) of -3.1% from 2017 to 2024 as cheaper SD-WAN solutions replace more expensive branch office routers[1].

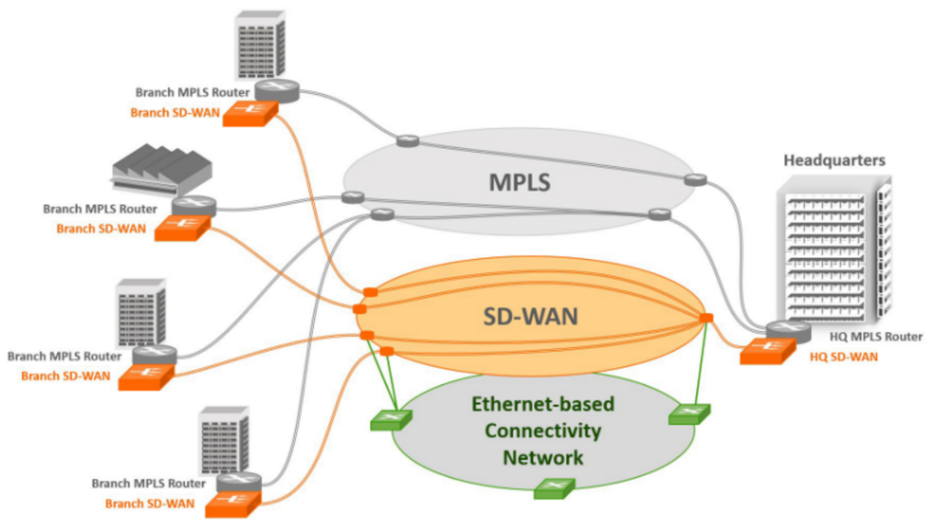

Figure 1: SD-WAN as backup to MPLS WAN connectivity

Nevertheless, one should not make the mistake of thinking that SD-WAN is less complex or intelligent. Quite the contrary in fact. The latest generation of SD-WAN solutions are expected to be application-aware and capable of determining the optimal path through the WAN in real-time on a per application basis. This includes cloud services with direct offload from the branch to co-located cloud edge services. In addition, sophisticated security capabilities are provided to support Zero Trust Network (ZTNA) and/or Secure Access Service Edge (SASE) concepts.

The sophistication of modern SD-WAN solutions allows a true de-coupling of WAN connectivity from the underlying transport mechanism, which is predominantly Ethernet based. It also means that there is little insight into the data transport layer, only to the

[1] Source: Gartner Magic Quadrant for WAN Edge Infrastructure

tunneled SD-WAN connections that are supported. As data consumption continues to grow and more sophisticated and demanding Internet-based services contend for public Internet resources, the underlying data transport layer networks supporting SD-WAN are becoming more dynamic and unpredictable. Multiple choices can be available, such as fixed broadband access over copper or fiber, Fixed Wireless Access (FWA) or 5G mobile broadband, but when should one choose these options, and can one be sure that the current measured performance will be maintained?

Testing sophisticated SD-WAN appliances, virtual functions and services becomes challenging when the SD-WAN solution is intelligent, and the underlying data transport connectivity is dynamic and unpredictable. To ensure that SD-WAN solutions are resilient and robust in the face of unpredictable Ethernet-based network behavior, it is necessary to emulate potential network issues and verify that performance measurement capabilities and SD-WAN policies react appropriately.

Ethernet impairment testing can be used to emulate the underlying Ethernet-based data transport network and introduce errors, otherwise known as impairments, such as packet loss, latency, jitter and packet re-ordering that are likely to occur in dynamic public Internet networks. This technique can be used to create a “performance landscape” where the contours of performance limits can be plotted and understood.

Plotting the SD-WAN performance landscape can help improve the algorithms supporting intelligence in SD-WAN edge appliances and virtual functions as well as the policies and orchestration solutions used to control SD-WAN networks. It can form the basis for Service Level Agreements (SLAs) and it can be used to reproduce WAN connectivity issues that are difficult to resolve.

In short, plotting the performance landscape provides reassurance that the SD-WAN solution will work, no matter what.

WHAT IS SD-WAN?

Software Defined Wide Area Network (SD-WAN) is a centrally controlled wide area network based on virtual connections across different connectivity networks. Inspired by Software Defined Networking (SDN), SD-WAN controllers define routing paths through the connectivity networks that meet performance requirements and other policies. This simplifies the management of WANs across multiple connectivity options while enabling fast, secure establishment of connectivity to new locations.

MEF has defined the components of an SD-WAN service in the MEF 70 specification (see MEF 70 Letter Ballot - SD-WAN Service Attributes and Services):

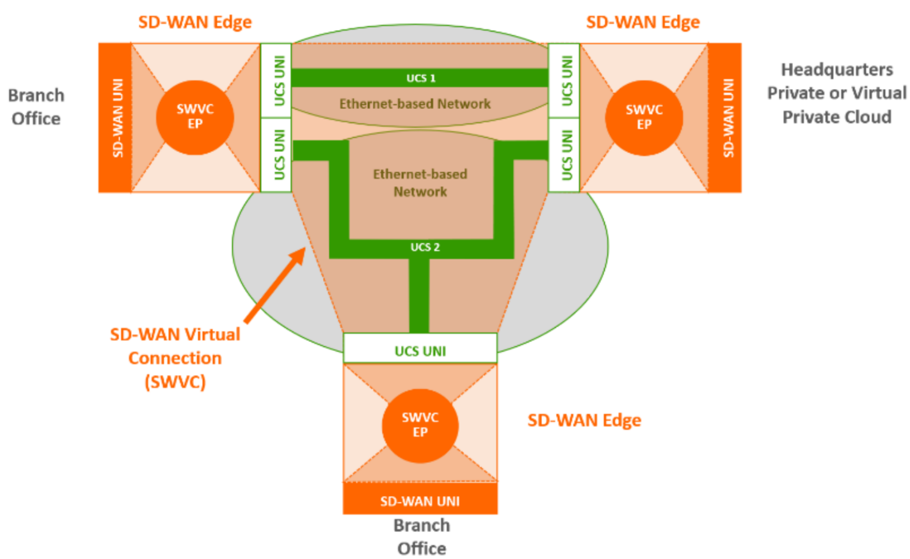

Figure 2: How SD-WAN's work

The SD-WAN is composed of several SD-WAN virtual connections (SWVC) that transport WAN traffic over a specified route in the Underlay Connectivity Service (UCS), which is often an Ethernet-based connectivity network. The SD-WAN Edge provides access to the SD-WAN and includes the SD-WAN User Network Interface (UNI), which is the demarcation point for SD-WAN service provider. It also includes SWVC End-Point (EP), which includes all the functions required to be performed on ingress and egress packets to support the SD-WAN service, such as measuring performance, imposing policy and routing packets to the appropriate SWVC. Finally, the SD-WAN Edge includes the UCS UNI, which connects to the underlying Ethernet-based connectivity network and establishes the SWVC.

For many enterprises, the SD-WAN service is provided by an internal IT organization, but a growing number of enterprises are sourcing SD-WAN services from SD-WAN service providers. SD-WAN 3.0 is a reference to 3rd generation SD-WAN services where the SD-WAN connectivity is provided over a private, managed underlay connectivity network owned by the service provider and is the latest evolution of SD-WAN.

SD-WAN SOLUTIONS HAVE EVOLVED

SD-WAN has evolved dramatically over the last decade. SD-WAN was originally introduced as a public Internet-based backup to expensive MPLS connections and later evolved to provide augmented bandwidth where non-critical applications could be transported over public Internet connections.

At the time, public Internet connections were not deemed reliable enough to support business critical applications, like voice services. However, with the mass adoption of cloud services and video streaming, public Internet services have become more robust, and many enterprises question whether expensive MPLS connections and branch office routers are even required.

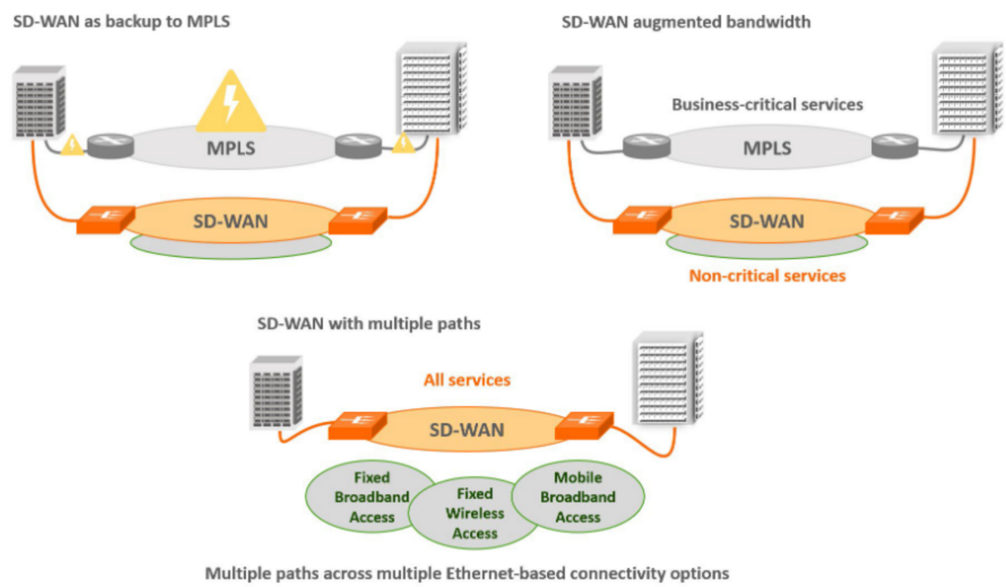

Figure 3: Evolution of SD-WAN

SD-WAN solutions including equipment, virtual functions and controllers, have also evolved in line with these developments. As more services move to the cloud, application performance requirements have become more differentiated and demanding. The ability to select the right SD-WAN path with the appropriate performance profile for a given application is now a basic requirement.

According to Gartner[2], a modern SD-WAN solution should support the following capabilities:

- Routing
- Application recognition
- Path selection
- Virtual private network (VPN) and Layer 4 firewall
- Native advanced security
- Cloud gateways
- Application performance optimization

Application performance optimization, in particular, is an important motivation for modern enterprises who increasingly rely on cloud services and applications rather than traditional on-premise software installations. SD-WAN solutions are expected to recognize applications, understand their requirements and prioritize available paths through the SD-WAN. To do this, the SD-WAN must be able to accurately determine the current status and performance of available SD-WAN paths and their supporting connectivity links in real-time.

MODERN SD-WAN FOCUS ON APPLICATION QOE

Optimizing the Quality of Experience (QoE) of enterprise applications is one of the main motivations for investing in modern SD-WAN solutions. A number of capabilities and features are provided in SD-WAN solutions for detecting and reacting in real-time to WAN issues so that application delivery is not affected.

Application awareness and path performance

The first step is recognizing, classifying, and prioritizing individual applications. This is achieved with advanced Deep Packet Inspection (DPI) capabilities that examine packet flows in real-time and identify the applications supported by each packet flow. The list of classified and prioritized applications enables SD-WAN policies to be created that will ensure that each application is transported on SD-WAN paths that can meet specified performance requirements.

To select an appropriate SD-WAN path, it must be possible to determine the performance of each SD-WAN path. Each SD-WAN endpoint uses mechanisms like

[2] Source: Gartner Magic Quadrant for WAN Edge Infrastructure

Bidirectional Forwarding Detection (BFD) to determine if the SD-WAN path is available and functional. BFD sends “Hello” packets across the SD-WAN path to neighboring nodes and if a reply is not received within a given timeframe, then the path is determined to be down. The same BFD “Hello” packets can be used to measure packet loss, latency and jitter on the SD-WAN link providing a reliable determination of the performance of the link at the time of measurement.

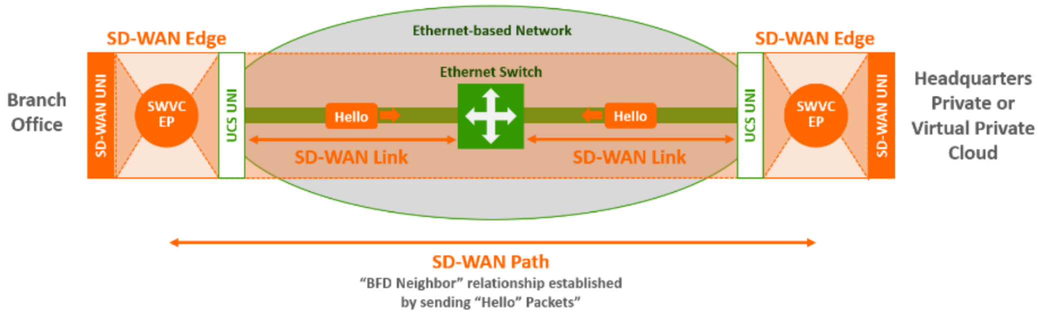

Figure 3: BFD "Hello" packets over SD-WAN

SD-WAN endpoints are therefore constantly monitoring application flows in real-time and making adjustments to compensate for network issues encountered. For example, SD-WANs can determine application, network and server response times and can understand application transaction response time variation on different SD-WAN paths.

TCP optimization

SD-WAN endpoints are also capable of optimizing application performance using TCP optimization. For example, it is possible to proactively discover the Maximum Transmission Unit (MTU) size on each SD-WAN link and report back to the application that the MTU has been exceeded. This allows the application to reduce the packet sizes used and therefore avoid fragmentation, which can cause issues in SD-WAN networks. In addition, TCP optimization techniques, such as TCP Selective Acknowledgement (SACK) can be used to prevent unnecessary re-transmissions and large initial TCP window sizes to maximize throughput and improve QoE.

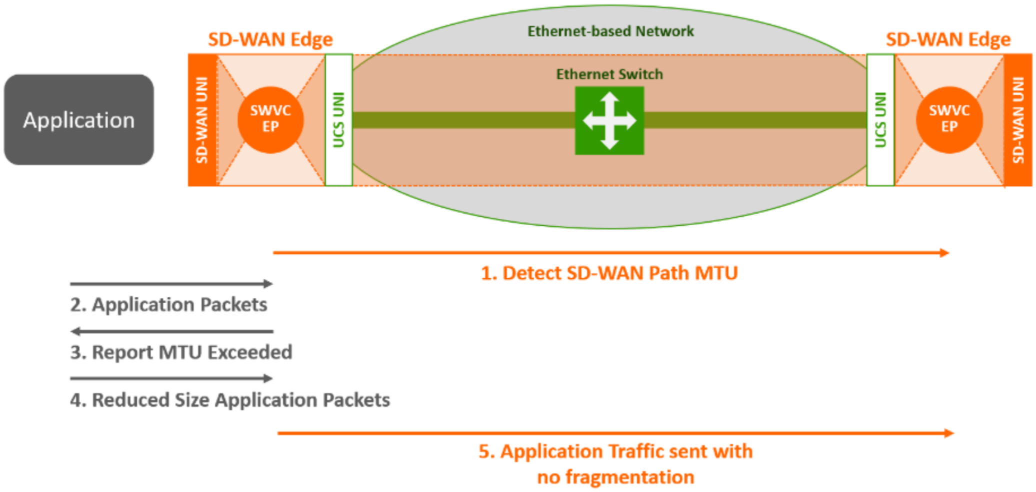

Figure 4: PMTUD Mechanism

Modern SD-WAN solutions are therefore quite sophisticated and intelligent with the capability to react in real-time to network issues while optimizing application QoE.

TESTING MODERN SD-WAN PERFORMANCE

Modern SD-WANs are more sophisticated and intelligent and increasingly rely on more dynamic public Internet connectivity options. This poses new test challenges.

In more traditional WAN architectures, the links connecting branch offices to headquarters are engineered with relatively predictable behavior. If an MPLS link is engineered to provide a certain bandwidth and performance characteristics, then one can safely assume during testing that the MPLS link will display these characteristics in normal operation. In other words, the WAN is relatively static and predictable.

In modern SD-WAN networks, however, this can no longer be assumed. Public Internet connections are more reliable and resilient than in the past, but this is a public network. Subscribers must contend for bandwidth traffic with other subscribers and their applications. Since there are an increasing number of demanding services that now rely on the Internet, it is difficult to predict traffic patterns and usage behavior. Congestion can occur leading to packet loss and buffering, which leads to latency and jitter issues. In short, the SD-WAN is now more dynamic and less predictable.

This is one of the reasons why modern SD-WAN endpoints have become more intelligent and capable of real-time performance measurement that can react to any congestion or availability issues in the SD-WAN. But how can this capability be effectively tested? How can SD-WAN vendors ensure that their SD-WAN solution will behave appropriately in the face of SD-WAN congestion and outage issues? How can SD-WAN service providers be sure that the sourced SD-WAN endpoints can meet their performance requirements? How can enterprises be sure that the SD-WAN service and solutions they have purchased can meet their SLA requirements?

Typical SD-WAN issues

Some issues can have outsize impact on application QoE. These include:

- **Poor quality links:** when SD-WAN links have poor quality, this can lead to packet loss. A packet loss of only 0.1% can reduce throughput by as much as 90%[3]. Poor quality links can result in packet corruption, which can mean that the integrity of the packets received is compromised and the packets must be discarded and retransmitted.

- **Network congestion:** when the routers and switches in the SD-WAN path are congested, packets can be dropped or buffered leading to latency and jitter. For many applications, this is not an issue, but for time-sensitive applications like voice, video and business-critical applications, latency and jitter can have a serious effect on the QoE for the user.

- **Mismatched MTUs:** The Maximum Transmission Unit (MTU) defines the maximum size of packet that can be transmitted. If the application uses an MTU that is larger than devices in the network, then the packet must be fragmented. The loss of a single fragment means that all fragments need to be resent leading to major latency. Fragments can also be received out-of-order requiring reordering, which also introduces delays.

- **Multi-path transmission:** to ensure delivery of packets, some SD-WAN implementations will send the same packets on multiple paths. This will result in duplicate packets being received at the SD-WAN endpoint that must be discarded requiring more processing as well as consuming additional network bandwidth that can cause congestion. It can also result in out-of-order packets that need to be reordered, which introduces delays. It should be noted that multi-path transmission is often chosen to ensure delivery of sensitive packets, but can be avoided as it is a chosen behavior.

Typical mitigation strategies

The following are typical mitigation capabilities in SD-WAN endpoints to avoid the issues listed above.

[3] Source: An Introduction to High Performance Network Connectivity | Fibre Systems (fibre-systems.com)

- **FEC for packet integrity:** to recover packets that can get corrupted, SD-WAN endpoints can use Forward Error Correction (FEC) methodologies. The transmitting SD-WAN endpoint can insert a parity packet for every four data packets allowing the receiving SD-WAN endpoint to reconstruct a lost packet based on the parity value.

- **Application traffic prioritization:** to ensure that latency-sensitive packets are not sent on SD-WAN links that can be prone to congestion, applications are prioritized and matched with high-quality links and immediately re-routed should the performance on the link come below a specified threshold.

- **Path MTU Discovery:** to prevent packet fragmentation, SD-WAN endpoints can use Path MTU Discovery (PMTUD) to discover the supported MTU on SD-WAN links and report back to the application. This enables the application to adjust the MTU accordingly and avoid packet fragmentation.

PLOTTING THE PERFORMANCE LANDSCAPE WITH IMPAIRMENT TESTING

Plotting the performance landscape means contouring the limits of network performance. It can be considered a stress test where network conditions are emulated to see how the network functions and equipment - as well as the services they support - respond to different events. The best approach for achieving this is to use Ethernet impairment testing.

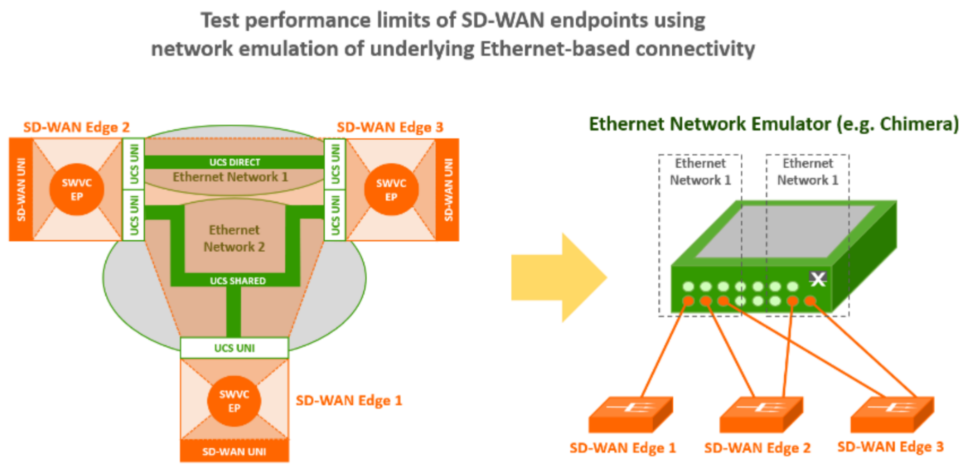

Figure 5: Emulation of SD-WAN Ethernet connectivity

A network emulator is a device that emulates network behavior so that it appears to attached devices as if they are connected to a network. With a network emulator, it is possible to change how the "network" is behaving using configurable parameters. It is also possible to introduce error conditions or impairments, such as high latency, jitter, packet loss or other errors that emulate what could happen in real networks.

What is impairment testing?

In Ethernet packet networks, issues can occur that hinder the timely and orderly delivery of data. These are "impairments" and can be issues like packet drops, packet duplication, packet latency and jitter or packets arriving out-of-order.

In real-world packet networks like Ethernet, impairments can occur under network congestion or network switch overload. The issues occur spontaneously, often in very short periods of time, but can significantly impact network performance and service delivery.

To test the resilience of network equipment in such situations, a mechanism is needed to introduce impairments into the network. However, this can be difficult to control in a live network. A network emulator simulates the behavior of an Ethernet network, but since it is a self-contained appliance, it provides predictable performance without any outside influences. This allows impairments to be introduced in a controlled manner and their effects on devices under test to be examined.

The impairments can be introduced individually as a single event or based on a schedule and distribution profile that emulates how a specific Ethernet network behaves. For example, introducing a packet drop impairment at regular intervals over a specific period of time and for specific types of Ethernet traffic. Impairments can also be combined to emulate complex error conditions. This can be used to reproduce behavior observed in real networks, but which is hard to isolate in a live network.

Ethernet impairment testing is a cost-effective alternative to field testing. Tests can be performed in the lab at fraction of the cost. In addition, there are no limits to the range of tests that can be performed including the ability to automate allowing continuous repeated testing. This allows specific scenarios to be tested that can be difficult to test in the field.

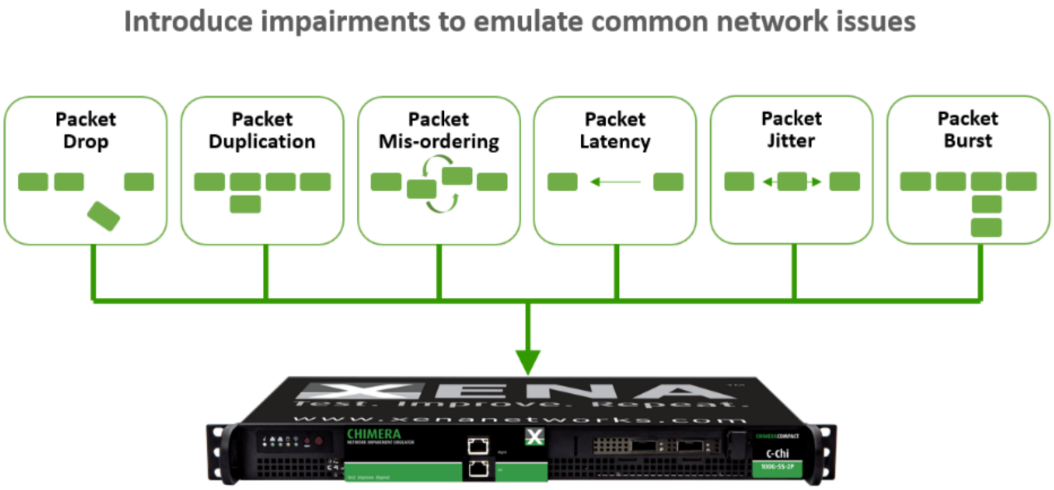

Figure 6: Network emulator impairments

Using a network emulator and impairment testing, it is possible to introduce error conditions in the emulated underlying Ethernet network supporting SD-WAN paths and links. This will allow the behavior of SD-WAN endpoints as well as SD-WAN controller policies to be tested under severe impairment conditions. By plotting the results for different types of impairment test, it is possible to develop an overview, or landscape, of the performance limits of the SD-WAN endpoints and controllers under test.

KEY AREAS FOR SD-WAN IMPAIRMENT TESTING

There are key areas of SD-WAN operation that need to be tested to ensure that SD-WAN solutions react appropriately to issues on a given SD-WAN connection:

- Reaction to packet loss
- Reaction to latency and jitter
- Reaction to out-of-order packets
- Reaction to duplicated packets

Packet loss

By introducing packet loss impairment, it is possible to emulate SD-WAN connections with degraded quality. It tests the detection and reaction capabilities of SD-WAN endpoints at different packet loss rates and can be used to determine the packet loss performance limit of the SD-WAN solution.

Latency and jitter

By introducing latency and jitter impairment, it is possible to emulate network congestion, queueing and buffering issues in the SD-WAN connection. The detection and reaction capabilities of the SD-WAN endpoints can be determined at different rates of impairment to determine the latency and jitter performance limits of the SD-WAN solution.

Out-of-order packets

By re-ordering the packets, it is possible to emulate the effect of packet re-routing and the delay introduced for packet re-ordering on an SD-WAN connection. This tests the ability of the SD-WAN endpoint and SD-WAN controller to detect and react to potential routing and re-routing issues.

Duplicated packets

By duplicating packets, it is possible to see the impact of introducing multiple paths for more secure delivery of critical applications. The SD-WAN endpoints need to identify and discard duplicates, which can have an impact on overall performance of the SD-WAN solution.

Combining and automating scheduled tests

The individual tests provide valuable insights but combining and scheduling these tests based on configured distributions enables realistic SD-WAN behavior to be emulated. This provides a more realistic insight into how the SD-WAN solution will behave under extreme conditions.

IMPAIRMENT TESTING ADDS CRITICAL DIMENSION TO PERFORMANCE ASSURANCE

Ethernet impairment testing adds a new, critical test dimension that provides invaluable insight and tangible benefits.

Manage dynamic network behavior

As underlying Ethernet-based connectivity networks become more dynamic and need to support a broader range of disparate applications, SD-WAN solutions need to be able to react to unforeseen changes in the network in real-time and ensure that application QoE is not affected. Ethernet impairment testing can be used to stress-test SD-WAN solutions before they are deployed to ensure that they can react appropriately.

Cost-effective alternative to field-testing

Plotting the performance landscape is not new. However, it is usually performed in the real-world during field trials and live deployments where feedback is used to improve new versions of SD-WAN solutions. This is costly, risky and takes time.

Ethernet impairment testing can be performed in a lab environment and address most potential conditions, events and other performance cases that could be encountered in live deployments. In fact, a meticulous impairment test can provide more accurate and complete information on performance under various conditions. Feedback from field trials and live deployments can be incomplete, as extreme conditions might never be encountered. This is a worst-case scenario, as failure when extreme conditions are encountered could have catastrophic consequences.

The ability to automate these tests with precise schedules and test distributions allows specific scenarios and conditions to be emulated over time that provides invaluable insight into the longer-term effects of Ethernet network configurations that can eventually lead to major issues.

Emulate observed conditions to understand causes

SD-WAN architectures can be extremely complex and when a performance issue is encountered, it can be difficult to determine the root-cause. Ethernet impairment testing can be used to reproduce the behavior observed and provide insight into what the root-cause could be. This has the advantage of not interfering with the existing network that is also supporting multiple other services that are not affected.

Plan new features and deployments

Network emulation allows new and more intelligent SD-WAN features to be tested accurately before deployment as the existing network conditions can be emulated accurately. More sophisticated algorithms and policies can be tested and compared. Ethernet impairment testing can focus on finding the breaking point limits of new features to understand the limits within which performance can be assured.

USE CHIMERA TO PLOT THE SD-WAN PERFORMANCE LANDSCAPE

Ethernet impairment testing can be used to plot the performance landscape of SD-WAN solutions.

TEST MODULE (Chi-100G-5S-2P)

Chimera is currently available as a 2-slot test module for a ValkyrieBay chassis, the Val-C12-2400G, where ValkyrieManager is used to emulate impairment as part of it standard user-interface.

STANDALONE CHASSIS (C-Chi-100G-5S-2P):

ChimeraCompact is a standalone version of Chimera where the test module is installed in a quiet desktop-sized chassis.

Xena's Ethernet impairment testing product is named Chimera.

Chimera can be used to plot the performance landscape of SD-WAN solutions by emulating the underlying Ethernet-based connectivity networks and issues, such as packet drops, latency and jitter. This provides insight into the resilience and robustness of SD-WAN solutions in the face of real-world Ethernet network issues.

Controlled testing at scale

Chimera supports multiple ports of 10G, 25G, 40G, 50G and 100G allowing a number of devices under test simultaneously. A comprehensive range of Ethernet impairments can be introduced on a per port and per-flow basis. Each impairment can be configured for a specific behavior in a specific period that can be repeated. This allows complex combinations of tests that can accurately even the most complex network behavior. The network emulator provides an easy-to-use and understand Graphical User Interface (GUI). Nevertheless, the solution is entirely script driven enabling configuration and reporting to be automated and controlled using Chimera's strong scripting interface.

Ideal for SD-WAN connectivity emulation

Chimera is fully integrated with the Xena Networks Valkyrie stateless traffic generator allowing a complete testing solution for SD-WAN connectivity emulation. Chimera can be configured to emulate the Ethernet connectivity issues that can affect SD-WAN performance and selection of the most optimal SD-WAN connectivity option. Custom Ethernet impairment distributions allow tailored impairments to be defined that emulate specific Ethernet network behavior all of which can be automated with configurable durations and repetition.

For more information see https://xenanetworks.com/chimera/.

>> Book a consultation with a Xena tech expert