



SD-WAN Benchmarking Best Practices for QoS

by The Tolly Group

The goal of this paper is to outline considerations and methods for empirical evaluation.

Presented courtesy of

infovista
KNOW YOUR NETWORK™

Executive Summary

While Internet bandwidth increases continually for remote sites, the application demand for that bandwidth can almost be guaranteed to grow at an even faster rate. Uncontrolled application contention for that limited resource results in degraded quality for applications. Quality of Service, or QoS, is the broad-brush term for technology that resolves the application performance problem. QoS solutions prioritize applications and redistribute available bandwidth so that important applications receive sufficient bandwidth to perform as desired. In this report, Tolly will outline key concepts and considerations for benchmarking and validating the capabilities of QoS solutions.

There are dozens of vendors offering SD-WAN solutions and QoS is likely to be an important part of every solution. QoS is implemented many different ways. This is not surprising since there is no industry-standard way of implementing QoS. While there are standard ways of marking the desired QoS level of packets, that is where it ends. The implementation is up to the vendor. Period. The number of queues, prioritization algorithms, transport health tracking, etc. will differ across SD-WAN vendors.

At the most general level, QoS involves imposing order on the chaos of networks which, without outside influence, work on a first-come, first-served basis. The mechanisms for controlling traffic will differ in their granularity and perhaps in effectiveness when it comes to providing QoS to YOUR applications.

For example, virtually every solution that you will encounter will be able to prioritize voice traffic over, say, file backups. More granular and more sophisticated QoS will be needed if, for example, you have “sub-categories” of apps like encrypted voice or SaaS apps that terminate on a public cloud.

Your application load is dynamic and, often, so is the WAN - especially if you use Internet links. Always consider how your solution can adjust dynamically to changing conditions.

The best practice is to benchmark what is important to YOUR company requirements. Ultimately, that is what really matters. In these pages, Tolly will outline how SD-WAN solutions can be evaluated empirically before deployment so that the most appropriate solution can be selected and post-deployment “surprises” can be avoided.



Contents

- Concept-in-a-Nutshell*
- Scope*
- Business Goals*
- Benchmarking Goals*
- No Standard Tests*
- Lab/Environment Requirements*
- WAN: Simulated or Real*
- WAN: Static or Dynamic Bandwidth*
- Benchmarking Tools: Commercial*
- Benchmarking Tools: Open Source*
- Performance Is Relative*
- “Before-After” Benchmarking*
- Test Metrics*
- WAN Configuration: Key Attributes/Impairments*
- WAN: Other Considerations*
- WAN: MPLS Queues Shortcoming*
- Test Preparation*
- QoS Configuration*
- SD-WAN Management Analytics*
- Caveats*
- Test Group #1: QoS By Application*
- Test Group #2: QoS With Three or More Applications*
- Test Group #3: QoS Within Applications*
- Test Variations: QoS With Multiple WAN Links*
- Test Group #4: QoS - Send Rate From Many To One*
- Advanced Topics*
- Next Steps/Follow-On Work*

Concept-in-a-Nutshell

WANs are often overwhelmed with traffic. Intelligence is needed to prioritize and control the flow of traffic. QoS is the umbrella term for that intelligence.

Scope

This document aims to provide practical, strategic guidance to enable users to design benchmark scenarios for QoS solutions.

This document is not a step-by-step cookbook for running tests. Tolly provides consulting services that can assist organizations in realizing specific tests.

Infovista Insights...

As more customers demand cloud-based applications, such as, Office365, Skype for Business and Teams, enterprises need to have a system that enhances and provides full control over the user experience. This is critical for customer retention and employee productivity.

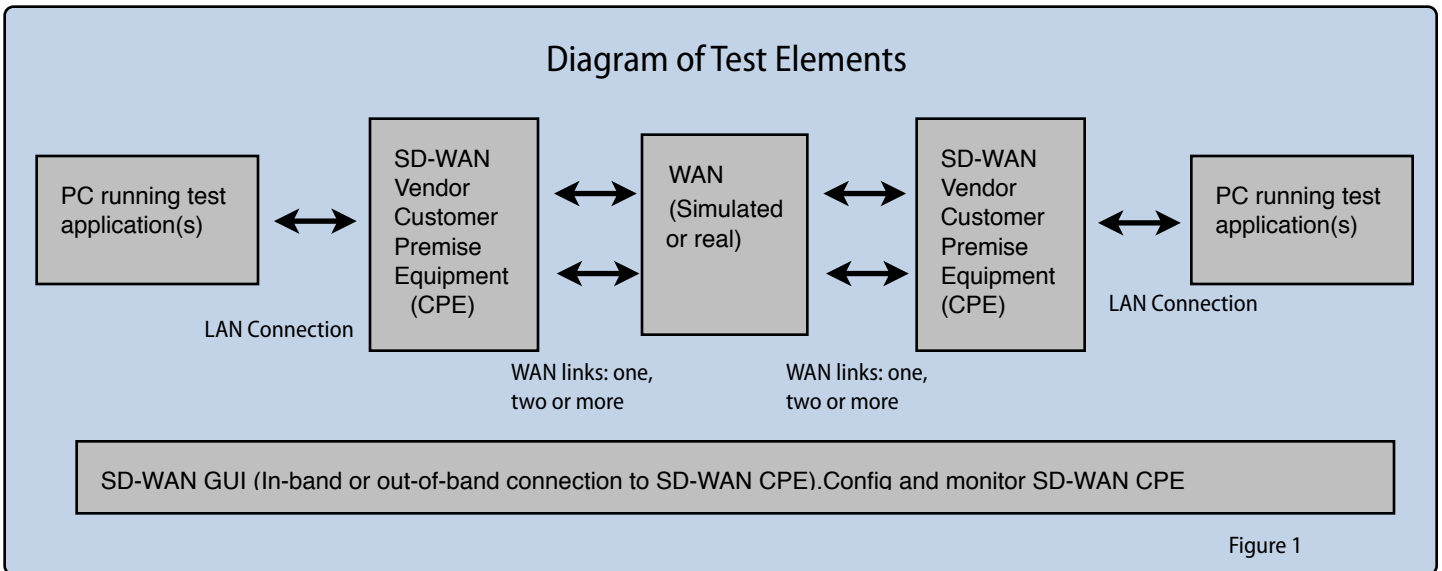
Infovista provides Application Intelligence+ in our Ipanema SD-WAN solution for managing application performance. It prioritizes bandwidth usage for your most critical applications via dynamic class-based QoS and routes application traffic to the best path for each individual application user session, all in real-time. Application Intelligence+ works in the background to maximize the return of your SD-WAN investments, by allocating resources autonomously, automatically and dynamically for high delay and jitter sensitive cloud applications while optimizing transactional applications and high data transfer applications in parallel.

Application Intelligence+ ensures that your critical business applications are always-on regardless of network conditions, working nonstop by dynamically adjusting all application traffic flows, per session, based on real-time network conditions on all network links. For every critical application, it ensures that the performance and Quality of Experience (QoE) never dips, even during peak hours.

With better operational control over your SD-WAN, you can add and move branch sites faster and quickly manage changes to existing configurations. With cloud-based orchestration, Zero-touch provisioning and automation, enterprise operations can be further simplified with Ipanema SD-WAN.

In addition, Ipanema’s Application Intelligence+ provides the tools to maximize customer experience: You can now contextualize and personalize offers and services using insights to proactively identify opportunities to upsell. This has a direct, measurable impact on customer satisfaction levels and net promoter scores as well as customer sales conversion rates thanks to the consistency in performance delivered to the user experience of your critical business applications.

Finally, due to the advanced class-based dynamic queueing in Ipanema’s QoS coupled with dynamic path selection, Application Intelligence+ enables the highest utilization rate of every network link – often reaching 95% utilization rates without suffering from packet loss or high error rates. It tunes the network automatically and continuously to deliver more for less, saving millions of enterprise dollars in optimized bandwidth without the need for added bandwidth. Unlike solutions that rely on FEC techniques to address application performance by using increased bandwidth for duplicate data, Ipanema SD-WAN provides optimal performance for your critical business applications while maintaining the most efficient use of network resources resulting in higher employee productivity, increased customer satisfaction, greater revenues, and lower operational costs.



Business Goals

Before any discussion of benchmarking specifics the team should spend the required time to understand the business priority of each application. Only with that knowledge can one effectively design a relevant and meaningful set of test scenarios.

Benchmarking Goals

Outline increasingly granular tests modeling common WAN transport behaviors that will illustrate specific capabilities to provide QoS to applications.

No Standard Tests

While several IETF and IEEE standards address how QoS priorities are indicated in frames and packets, there are no generally recognized de facto "standard" tests that can be used to benchmark QoS.

Lab/Environment Requirements

Many QoS tests can be performed using a very basic set of software tools that can be set up quickly and with minimal cost.

Minimum lab requirements: 1) SD-WAN vendor devices (solution under test) 2) Wide-area network (WAN) emulation/simulation system, 3) pair of client computers that will communicate across the simulated or actual SD-WAN environment (at least one on either side of the WAN), and 4) one or more applications that will be used to benchmark QoS effectiveness. See Figure 1.

Specifics will be addressed later in this paper.

WAN: Simulated or Real

Whenever "WAN" is referenced in this document, it is safe to assume that we are discussing "simulated" or "emulated" WANs. These two

words are generally used interchangeably. (Note: "simulated" WANs are different in that mathematics are used to model the WAN. But for the sake of this discussion, they are the same.)

WAN: Static or Dynamic Bandwidth

Traditional private, leased-lines are static. As noted elsewhere in this document, broadband link speeds will vary from the nominal level. Furthermore, many SD-WAN implementations will use a combination of both static and dynamic links. It is important to keep this in mind and attempt to match your actual environment with whatever simulated or actual WAN links that you will use for your benchmarking.



Benchmarking Tools: Commercial

WAN Environment:

There are various vendors of commercial test tools that can emulate the WAN environment.

Apposite Technologies offers two lines: 1) Linktropy WAN Emulators (up to 1Gbps throughput) and, 2) Netropy Network Emulators (up to 40Gbps). Ixia offers the Network Emulator II product

Virtually every SD-WAN solution provides value and QoS by using two WAN links (or possibly more) between its CPE and the WAN. Thus, simultaneous support for at least two WAN links is an essential requirement for your test environment.

Application Environment:

To simulate more complex applications such as voice over IP (VoIP), a commercial application such as Ixia's IxChariot would prove useful.

The primary reasons for using a simulated application are that such applications integrate measurement tools and provide test automation.

PC-based tools such as IxChariot can be used to send multiple different traffic types simultaneously. They can also be scaled by adding more PCs ("endpoints in Ixia parlance") into the environment.

For users requiring very high traffic load, purpose-built test tools are available that can generate L4-7 traffic.

Ixia offers its IxLoad application that runs on its hardware platform or is available in a virtual edition. Xena Networks also offers a hardware-based L4-7 solution.

Benchmarking Tools: Open Source

Open source solutions have the benefit of being free to use but typically lack any formal support and typically will have more limited functionality than the commercial alternatives.

WAN Environment:

These solutions will typically be deployed on a standard Microsoft or Linux machine outfitted with two LAN interfaces. Each interface represents an ingress/egress point of the WAN.

Be sure to use machines that have sufficient CPU power, RAM and network interfaces to run the WAN environment that you wish to emulate.

WANEm is a commonly used open-source emulator. The Integrated Multiprotocol Network Emulator/Simulator (IMUNES) is another free tool based on FreeBSD and Linux.

Tolly has not worked extensively with these Open Source alternatives but it appears that they do NOT offer emulation of multiple WAN links simultaneously. Thus, one might require multiple

About The Tolly Group

IT experts with over 30 years of experience
We provide product benchmarking and analyst services to the end-user and vendor community.
info@tolly.com www.tolly.com

WAN emulators in order to simulate multiple WAN links.

Application Environment

One option is simply to use test versions of your actual business applications. (We say "test" because you obviously don't want to be updating production data in any target databases during testing.)

If it is not practical to use actual business applications, then it is important to choose applications that have similar characteristics to your actual applications.

The application characteristics of interest are typically: throughput, latency and jitter. (Latency and jitter are the most important attributes for VoIP sessions.)

At a more granular level, important characteristics include: packet size, packet frequency, session duration and TCP port.

Performance Is Relative

For QoS, there are no absolute metrics such as packets per second or percentage of theoretical throughput as you would use when, say, benchmarking LAN switches.

QoS benchmarking results are relative. That is to say, the



benchmarking generally is concerned with the performance of applications “before” and “after” the QoS solution is implemented.

“Before-After” Benchmarking

For the “before” scenario, the WAN is deliberately overloaded (“oversubscribed”). Application demand for bandwidth exceeds available bandwidth. This scenario can be created easily, typically by simply running a file transfer across the WAN. Depending upon your application environment, you may wish to run bi-directional file transfer in order to create congestion on both links.

The WAN interfaces will simply try to pass whatever traffic is presented. Essentially, it is “first come, first served.” The file transfer will typically send a stream of large packets to the receiving station. Depending, of course, on your WAN bandwidth, that single traffic stream could be sufficient to take all available bandwidth.

With the file transfer running, latency sensitive applications like VoIP or web transactions, which also send much smaller packets, are likely to spend time waiting in the queue and, as a result, deliver poor performance/quality.

The QoS solution that will be tested will intervene (when enabled) and will inspect the incoming traffic and manage the queues and flows dynamically to match whatever requirements the customer has

requested. In general, important traffic, such as VoIP and web transactions, should be prioritized over file-transfer type “backup” and data mover applications.

Test Metrics

The particular metric of interest will depend upon the application.

Necessarily, one will be more interested in measuring the performance of the high-priority applications rather than the low priority applications.

First, let’s identify and briefly define the metrics and then look at them in the context of the application.

Throughput:

This is the measurement of bits transferred over time. Throughput is usually reported as bits per second (often kilobits or megabits depending upon the speed of then WAN) and averaged over time.

Throughput is also often referred to as “goodput.” In this case, the throughput is being measured at the application level. Goodput, thus, does not include network overhead in the calculation and, thus, will be a lower value than throughput at the network level.

This document will use the term throughput generically. You should apply that term as appropriate to where you are making your measurements.

Throughput can also be calculated - or at least estimated - by knowing

the size of a file being transferred and the duration of the transfer.

For example, if you use FTP to transfer a 10MB file, you are transferring roughly 80 million bits. (Don’t forget, that there is overhead in the packets.) If the task duration is 5 seconds, you, you divide 80 million by 5 and determine that your effective throughput was approximately 16Mbps.

Commercial testing applications, such as Ixia IxChariot, will automatically calculate throughput.

Delay: Latency & Jitter

At the risk of being redundant, delay (typically caused by congestion) is typically the root cause of poor performance for VoIP. Poor performance being defined as poor voice quality.

Latency measures the elapsed time between packets, jitter measures the consistency of the delay. An ideal result would be low average latency with very low jitter. That is to say, latency that is both low and consistent. These are network-level measurements.

The longer (higher) the latency and the greater (higher) the jitter, the worse the VoIP conversation will sound.

Transaction Time and Rate

Ultimately, the performance at the network level will impact the time required to complete transactions and the number of transactions



that can be completed per unit of time (usually, per second).

Measuring duration and transaction rate with and without congestion (discussed below) can be a very useful way of measuring the effectiveness of your SD-WAN solution.

Keep in mind that the definition of transaction is whatever you determine to be. A single, web inquiry that completes in less than a second could be considered one transaction or a large file transfer lasting many minutes could also be considered a single transaction. What is important is to be consistent in your comparisons of transactions - comparing like to like.

Voice Quality Measurements: MOS, PSQM, PESQ

Because it is difficult to map the impact of latency and jitter measurements on the quality of a call, a number of measurement algorithms have been developed over the years.

Some test systems will provide an estimated Mean Opinion Score (MOS) - an estimated quality score - by evaluating the latency and jitter measurements and calculate an approximate quality score.

PSQM and its successor, PESQ, use voice samples transmitted across the conversation to calculate voice quality.

It is beyond the scope of this paper to go into any details of these VoIP quality calculation methods.

Suffice to say that if you are concerned about quality for VoIP applications you should be familiar with these scales and seek test tools that will calculate one of these VoIP quality scores.

“Score” for Non-VoIP applications

Unfortunately, there is no equivalent to MOS scores when evaluating non-VoIP transactions.

With file transfer applications, one can gauge the throughput, as noted elsewhere, by comparing duration of the transfer and/or calculating effective throughput in Mbps.

For transactional traffic, there are several aspects to evaluate. Transactions per second (TPS) will be a good gauge of the workload being processed. Just as important, if the metric is available, is response time.

With response time, keep in mind that everything is relative. Relative, that is, to the application response time across a network that is not under load.

One expert notes that some transactions, credit card transactions for example, can require several seconds under normal circumstances. This is because of the multiple interactions between the credit card system and the host.

For your benchmarking, pick several suitable transactions and benchmark them for several runs (3 to 5 samples) over the network when it is not oversubscribed.

These response times can provide a baseline for comparison with the results delivered by the SD-WAN solution when the WAN is under load.

SD-WAN Management Statistics

An important aspect of your SD-WAN solution is its management and reporting capabilities.

As the SD-WAN solution will have visibility to your application traffic on a flow-by-flow basis, a well-implemented management capability can provide you with all the important information that you need on application performance.

QoS Precision

Depending upon the SD-WAN solution, you will have different options with respect to “precision.” That is to say, one solution might simply let you choose a priority level without allowing you to provide a more specific definition of QoS. Another might allow you to define a percentage of bandwidth given to a particular application or group. Another solution might allow you to specify actual bandwidth amounts to be allocated.

WAN Configuration: Key Attributes/Impairments

You will want to configure the emulated WAN environment to match your actual environment. What is less obvious is that there are aspects of your current (or planned) WAN environment that you may not know for certain.



Bandwidth:

This should be easy. Confirm the bandwidth configuration for your network. Even though you are connecting to the network emulator via 100Mbps or Gigabit Ethernet, the emulator will restrict the bandwidth to what you configure for your test.

Don't forget that most Internet connections are asymmetrical. That is, the uplink (upload) and downlink (download) speeds are different. The downlink speed will typically be the faster of the two.

Thus, when building your test environment be aware of this and configure your gear appropriately.

Note that with your test applications, it might not be possible to congest the WAN. If that is the case, you can deliberately lower the WAN bandwidth for your test to create congestion deliberately. After all, your primary goal is to evaluate the SD-WAN solution's response to oversubscription and congestion.

Conversely, you can also use this test environment to see if having higher WAN bandwidth available would improve application performance.

Actual Bandwidth Capacity

The actual bandwidth and capacity of a broadband circuit can vary significantly over time (hour-to-hour or day-to-day). This must be recognized and should be a factor in your testing.

For example, one expert noted to Tolly that he had witnessed a broadband circuit that was rated at 16Mbps by the service provider deliver in excess of 22Mbps average on one day and then only 8.5Mbps on another day.

Thus, it is important to distinguish between the ordered capacity and the delivered capacity (which, as noted will likely vary with broadband links). The ordered capacity is bandwidth level being sold by the provider, the delivered capacity is the actual bandwidth available to consume for data transport.

Impairments:

Your WAN/Internet provider won't like to talk about this topic, but the reality is that impairments are a part of life for your WAN.

For each of these impairments, try to determine your actual network characteristics (see below) and start there. It is also advisable to run additional testing where you dial in impairments that are somewhat worse than you expect. This way you can determine the point at which certain impairment levels might cause significant problems for your network performance.

Latency

There is delay inherent in packets traveling across the network. If you are not sure about the latency in your network, you can use a "ping" terminal command or a "SpeedTest" type of web application to give you some

insights into what latency value to use on your emulator.

Jitter

To recap what was noted earlier, jitter refers to differences in delay across packets. Low or nonexistent jitter is generally the better condition for your network.

Packet Loss

This is an important parameter. Usually expressed as a percentage, you can specify the number of packets that the emulator should discard. This simulates packets that are lost (usually because of congestion) in an actual network. Note that some emulators might have very simple algorithms (e.g. discard every 50th packet) where others might apply more sophisticated and realistic algorithms to generating the packet loss.

One expert notes that it could be useful to know if your SD-WAN solution can load balance traffic AWAY from links that are known to be "high loss" or of low quality.

Bandwidth

It is important to remember that bandwidth can also be an impairment. Specifically, the variation that can and will occur dynamically on broadband circuits presents challenges to the SD-WAN solution.

The solution must be able to detect these changes and respond appropriately to variations in the amount of consumable bandwidth available.



WAN: Other Considerations

Multi-link

All emulators, naturally, will be able to simulate a WAN consisting of a single link between the source and target WANs. It is not uncommon for providers to offer two links into the cloud. Depending upon the provider (and their software) the second link could be only for failover or it could be used to provide load balancing and QoS functions. For example, if one link became degraded or congested, traffic could be migrated to the secondary link.

To validate such scenarios, you would need to have WAN simulation tools that can simulate two or more links.

WAN: MPLS Queues Shortcoming

Users of MPLS queues might believe that this will solve QoS problems. A subject matter expert noted to Tolly that these queues simply configure bandwidth for a particularly tagged set of traffic and that there is no mechanism to prevent oversubscription of these queues. This would result in congestion and poor performance.

Test Preparation

To recap some of the above. You will need to set up the WAN environment with at least two interfaces. You will need to connect a LAN to each of those interfaces. These LANs will represent source

and target locations for the applications that must traverse the LAN. (Don't forget that if you are using cable modem/ADSL, the link speeds will be different in each direction.)

Finally you will have to have your applications (or Ixia IxChariot app simulator) in place. At a minimum, you will need a high-priority and a low-priority application. For low priority, something like a file transfer works well as it will tend to consume a lot of bandwidth and provide the WAN congestion that you need in order to evaluate the QoS solution

QoS Configuration

For your "before" tests, you will simply want to test either without your SD-WAN QoS solution in place. Or, if it is in place, make sure that the QoS functions are disabled.

Just as there is no standard way of implementing underlying QoS technology, there is no standard method for specifying how you want applications prioritized. There is, however, common QoS theory and elements that run through all of the SD-WAN solutions.

QoS solutions will common elements like priority queues or priority levels. The number of queues may vary - and that will impact the potential effectiveness of the QoS.

Each solution will also have to deal with bandwidth. That is, ultimately, the resource that must be

managed. Be aware whether you must statically configure inbound and outbound bandwidth for your interfaces. Dynamic bandwidth approaches are needed to respond to variations in consumable broadband capacity.

Each solution will present you with a different interface. You should inform the provider of the QoS solution of what your desired application priorities are and have the vendor provide the appropriate test configuration.

There may be additional options related to whether each application or application session receives a fixed amount of bandwidth or a relative amount of bandwidth.

You should spend some time exploring the QoS options provided as that can provide you with key insights as to the potential power of the underlying QoS technology.

SD-WAN Management Analytics

With our test scenarios, the SD-WAN is a "black box." That is, we see what is going in and what is going out. We can measure the various network characteristics. We might be able to surmise what is happening "inside" but, generally, we cannot.

Thus, we need to rely on the what the management analytics component of the SD-WAN solution can tell us. While this information is useful for our testing, it is even more important for your long-term



understanding and management of your SD-WAN environment.

For our testing and in actual production use, the management analytics should provide performance information about the various applications - the more granular, the better. Both real-time and historical data should be available.

Additionally, the analytics should provide information about the various WAN links that are being used for transport. Link speed, latency, packet loss and other impairment conditions should be reported by the SD-WAN analytics component.

Caveats

QoS Testing necessarily involves creating congestion in the wide-area network, thus it is not advisable to perform these benchmark tests on any production network.

Please note, however, that some SD-WAN service providers will have test networks available where such tests can be run without issue.

Test Group #1: QoS By Application

Always start with attempting to prioritize a single application. By definition, this application will be an important, therefore high-priority application.

For many companies, this will be VoIP. For others, it might be real-time, transaction-oriented traffic

related to sales or customer service.

The lower priority, background traffic could be file transfer or backup. It could be traditional file transfer program (FTP) streams or perhaps a backup or sync that targets a cloud-based server such as Dropbox.

While the background traffic might be intermittent in your actual network, you will want to be sure it is active during your test.

Typically it is best to start the background application first and be sure that it has a sufficiently large file (or set of files) to transfer that it will stay active during your run of your high priority application.

Sometimes the QoS solution will have a management console and statistics gathering that will make documenting “before” and “after” performance much easier. Check to see if this exists. Be sure to clear all counters before running.

After you run your basic “before/after” scenario, take a quick look at your results. While everything is still set up, you might want to change some parameters of the QoS system and/or try out some different applications so that you have multiple sets of applications and results to compare when you evaluate your results.

After each test it is a good idea just to make sure that the results make sense. It is typically a good idea to run each test scenario at least three times so that you confirm

that the results are consistent. If those results are not consistent, you might want to run additional iterations. Most importantly, you would want to understand why the results are not consistent as you want consistent behavior in your production network.

Test Group #2: QoS With Three or More Applications

It is most likely that your network will not be so simple as to have only “important” and “background” applications.

More likely, you will have more gradations. You might have “critical” traffic that has to get through no matter what. This could be followed by one or more “important” applications. Further down the list you could have “background” applications that have to complete but without a specific deadline. Finally, you might have some “low priority” applications - say video streaming by employees- that are allowed when there is sufficient bandwidth but can simply be discarded when the network is busy.

Most SD-WAN QoS systems will provide for at least four classes or queues of priority. Work with the vendor to outline your priorities and have them provide appropriate system configurations.

This testing becomes more complex simply because one needs to orchestrate running four or more applications simultaneously.



Should you get to this level, you should seriously consider investing in a commercial-class application test tool that can provide automated application tests.

Run and evaluate as previously noted.

Test Group #3: QoS Within Applications

Until now, we have prioritized traffic “by application” but it is also quite possible that you will require prioritization within a single type of application.

For example, you might have to assure executives that their VoIP sessions will also have good quality but for other inter-office VoIP a lower quality is acceptable.

For such a situation the QoS solution must provide a way for you, the customer, to specify “more important” and “less important” VoIP. Perhaps it might be by IP address or IP subnet or QoS bit markings. In any case, you will need to work closely with your prospective SD-WAN provider to make sure that they are able to provide this granular level of QoS within a single application type.

Once configured, you can test this using the test methods explained above.

Test Variations: QoS With Multiple WAN Links

Until now, we have looked at the external aspects of QoS, that is, the applications being prioritized. Managing and optimizing WAN resources is an important aspect of SD-WAN solutions and an area where one might find differences across solutions.

These WAN link variations can be run with any or all of the other QoS tests in this document.

Blackout

This is the most straightforward scenario. Here, the active link is deliberately failed and the application performance monitored. Generally, the primary requirement will be that any application running across the failed link will be migrated to a secondary link without session interruption.

Brownout

In this scenario, the active link does not fail. Rather, it experiences degradation of some type that, correspondingly, degrades the performance of specified applications to an unacceptable degree. Upon reaching that threshold, the relevant application sessions are migrated to a secondary link.

The conditions (i.e. impairments) that trigger the session migration will depend upon both the capabilities of the SD-WAN solution and your business requirements.

First you need to evaluate the quality criteria that the SD-WAN solution monitors. This list could include: packet loss, latency, bandwidth usage. It could also extend to important “end-to-end” considerations such as transaction response time. As these criteria will likely vary across vendors, you are well advised to review these closely.

Then, you must match these up to your application’s business importance. VoIP is latency sensitive and will not provide usable conversations if latency exceeds certain levels. Thus, that could be a trigger for that application.

Single-Link Mode

While your planned deployment might (and should) only be with redundant links, it is important to understand the QoS behavior of your SD-WAN solution when only a single link is available.

Think of it this way: if the response of your SD-WAN solution to congestion is simply to offload to a secondary link, you are not getting an accurate benchmark of the queuing effectiveness. By removing the secondary link you can put a focus on the the single-link queuing capability of your solution. This worst case scenario can provide information on performance in network-under-stress situations that you will want to be aware of.



Transport Cost Awareness

It is most likely that not only will your various WAN links have different capacity, but different pricing as well. In some cases, the pricing could be usage based.

Thus, you should review any options that your SD-WAN solution offers with respect to awareness of cost. For example, can you specify that LTE cellular be used only in the event of a link failure and not just to distribute workload? The ability to configure your SD-WAN solution to be sensitive to relative cost of WAN links could be an important economic consideration.

Test Group #4: QoS - Send Rate From Many To One

Similar in many ways to head-of-line blocking in LAN switches, a test scenario that might be relevant for certain users is the “many to one.”

In this scenario, multiple senders are directing traffic to a single receiving node. In the process, there might be congestion resulting in the attempt to deliver the concentrated traffic.

This scenario will only be of interest should you have applications that direct traffic in this way.

Advanced Topics

WAN Efficiency

Some SD-WAN solutions will use a technique known as forward error correction (FEC) to provide QoS.

In this case, multiple copies of packets are forwarded across the network. This can dramatically reduce the impact of packet loss. The cost, however, is that WAN bandwidth is consumed by the additional, usually redundant packets. This is a trade-off that may or may not be acceptable to you. The important point is to be aware of it.

Next Steps/Follow-On Work

This document should assist in benchmarking core SD-WAN QoS and helping you build your short list of solutions. At that point, you should research whatever advanced QoS capabilities are available and identify which might be relevant in your environment. While you have your test environment operational, it would be a good time to configure some of those functions and determine if the vendor’s advanced features can provide measurable benefits to your application environment.

###



About Tolly...

The Tolly Group companies have been delivering world-class IT services for 30 years.

Tolly is a leading global provider of third-party validation services for vendors of IT products, components and services.

Tolly also assists medium-sized businesses and large enterprises evaluate, benchmark and select IT products for deployment.

You can reach the company by email at sales@tolly.com, or by telephone at +1 561.391.5610.

Visit Tolly on the Internet at:

<http://www.tolly.com>

Terms of Usage

This document is provided, free-of-charge, to help you understand whether a given product, technology or service merits additional investigation for your particular needs. Any decision to purchase a product must be based on your own assessment of suitability based on your needs. The document should never be used as a substitute for advice from a qualified IT or business professional. This evaluation was focused on illustrating specific features and/or performance of the product(s) and was conducted under controlled, laboratory conditions. Certain tests may have been tailored to reflect performance under ideal conditions; performance may vary under real-world conditions. Users should run tests based on their own real-world scenarios to validate performance for their own networks.

Reasonable efforts were made to ensure the accuracy of the data contained herein but errors and/or oversights can occur. The test/audit documented herein may also rely on various test tools the accuracy of which is beyond our control. Furthermore, the document relies on certain representations by the sponsor that are beyond our control to verify. Among these is that the software/hardware tested is production or production track and is, or will be, available in equivalent or better form to commercial customers. Accordingly, this document is provided "as is", and Tolly Enterprises, LLC (Tolly) gives no warranty, representation or undertaking, whether express or implied, and accepts no legal responsibility, whether direct or indirect, for the accuracy, completeness, usefulness or suitability of any information contained herein. By reviewing this document, you agree that your use of any information contained herein is at your own risk, and you accept all risks and responsibility for losses, damages, costs and other consequences resulting directly or indirectly from any information or material available on it. Tolly is not responsible for, and you agree to hold Tolly and its related affiliates harmless from any loss, harm, injury or damage resulting from or arising out of your use of or reliance on any of the information provided herein.

Tolly makes no claim as to whether any product or company described herein is suitable for investment. You should obtain your own independent professional advice, whether legal, accounting or otherwise, before proceeding with any investment or project related to any information, products or companies described herein. When foreign translations exist, the English document is considered authoritative. To assure accuracy, only use documents downloaded directly from Tolly.com.

No part of any document may be reproduced, in whole or in part, without the specific written permission of Tolly. All trademarks used in the document are owned by their respective owners. You agree not to use any trademark in or as the whole or part of your own trademarks in connection with any activities, products or services which are not ours, or in a manner which may be confusing, misleading or deceptive or in a manner that disparages us or our information, projects or developments.