SOLUTION BRIEF

Reflector-based Measurements

Zero-touch testing and monitoring for improved customer experience

Automating activation testing and service quality monitoring using reflection technology

Reflector-based measurements are ideal for quick, easy and cost-efficient implementation of both continuous monitoring and large-scale tests as they make use of embedded reflectors in existing network infrastructure, without the need for additional investment or configuration.

The Netrounds active testing and monitoring solution includes a variety of active *reflector-based* measurements, allowing you as an operator or service provider to significantly extend your testing capabilities and end-to-end network service quality visibility with ease. Place one Netrounds Test Agent at a central location in your network and bounce traffic off network devices, at the SD-WAN Edge for example, with reliable, standards-based reflector functionality. This allows you to test or actively monitor metrics such as throughput rate, packet loss and delay.

With this technique, you can take advantage of two key benefits:

1. Testing and monitoring from the end user perspective and location without the need to send technicians into the field
2. Installation of hardware or software at the customer site no longer required, decreasing both time and costs associated with testing and monitoring

Another key benefit of active testing and monitoring is the ability to be notified in real time of any service faults or SLA violations. Alarms and notifications allow network operations teams to immediately begin isolating issues through remote troubleshooting with Netrounds so that *problems can be fixed proactively*, before customers call in with complaints of service degradations.

Reflector-based measurements also assist in the troubleshooting phase of the network service lifecycle. Routers supporting reflector capabilities that are located between different measurement points can be used to further segment the network and isolate the problem remotely. This can be accomplished with our ICMP path trace feature, for example.

Further time and cost savings can be realized when automating reflection-based testing, such as service activation tests, ongoing monitoring and remote troubleshooting and segmentation through Netrounds' programmable APIs. APIs currently available include REST and NETCONF and YANG (see netrounds.com/resources for more on the APIs available).

Reflector-based measurements perform the following:

- Activation tests
- Ongoing monitoring
- Issue isolation and troubleshooting

Purpose of Reflector-based Measurements

Reflection-based measurements are carried out for three main purposes:

Activation testing: a test to validate a service before it is handed off to customers. Activation tests are intrusive by design as they should comprehensively stress-test the service with a larger traffic load and mix.

Ongoing monitoring: a virtually non-intrusive measurement that proactively validates that the service works continuously.

Troubleshooting: tests used to assist with issue isolation by increasing the information available for a service problem. The goal is to decrease mean time to repair (MTTR).

Deployment of a New Branch Office

A common use case for reflector-based measurements is testing of a new branch deployment, such as a new SD-WAN branch or a Business VPN connection. Traffic is sent from a Netrounds Test Agent, typically placed as a test head in the core part of the network or in the data center, to customer-premises equipment (CPE) such as routers located at the branch sites. These network elements reflect the traffic back to the Test Agent test head located in the operator core data center or enterprise headquarters data center.

As shown in the below diagram, where CPEs supporting standards-based reflection such as TWAMP, test traffic can be reflected from the Test Agent in the data center to the end points for measurements. One site has deployed a Test Agent at the end location, so traffic is sent from Test Agent to Test Agent, as well as into the Internet and the public cloud, for critical KPI measurements.

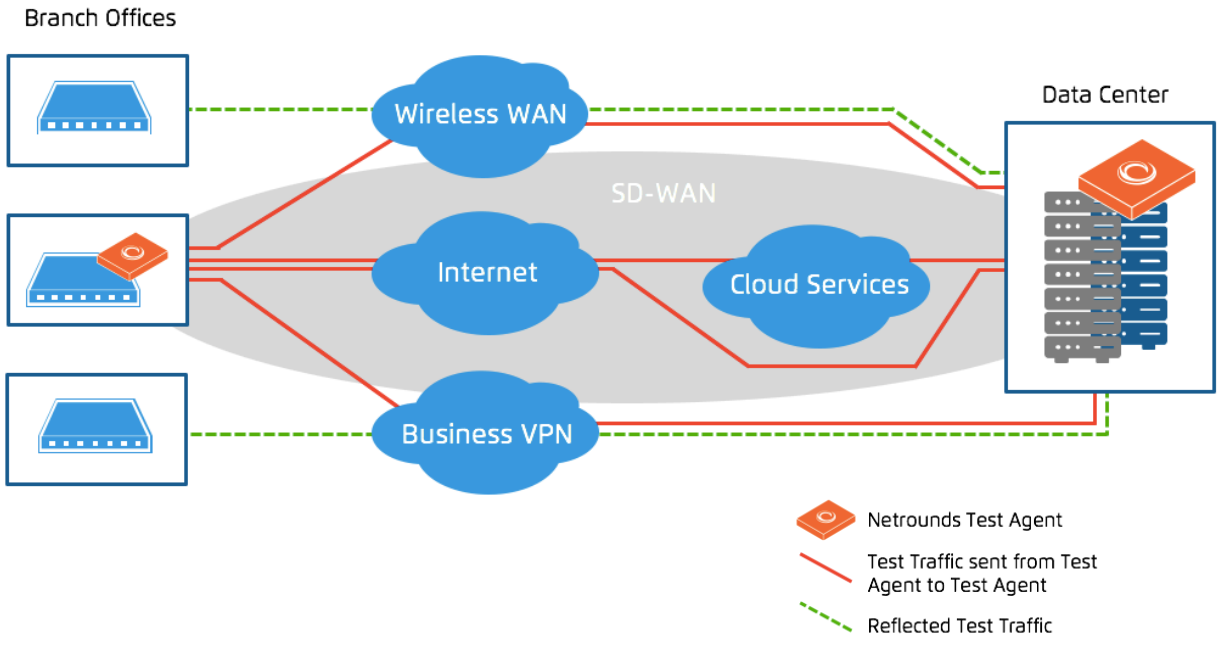

SD-WAN Architecture Image Credit: Juniper.net

It is important to note that other network topologies and Netrounds Test Agent locations are also possible. The above only serves as one example of a deployment model where numerous Test Agents and reflector-based measurements are used simultaneously with a multi-tenant Test Agent as the test head located in the data center.

Reflector-based Measurements Available in Netrounds

The following details the types of reflector-based measurements currently available in Netrounds.

ICMP Ping and UDP Echo

The most basic items in the toolbox of reflector-based tests are ICMP Ping and UDP Echo. A major advantage of ICMP Ping and UDP Echo is that support for these utilities is generally preconfigured in CPEs, so they are typically always available. The main application of these measurements is to test the reachability of hosts in an IP network. One item to note is that the maximum throughput of these tests may be limited, which is mainly due to network element limitations (see the "*High throughput limits: Software vs. hardware loopback*" chapter). If higher traffic rates are needed, UDP loopback may be used (see more on UDP loopback below).

Two-way Active Measurement Protocol (TWAMP)

Netrounds supports the use of TWAMP and TWAMP Light, both defined in IETF RFC 5357[1]. This protocol is widely available in many CPEs, and it can be used for measuring two-way metrics such as throughput, packet loss and delay, as well as one-way metrics if the reflector has sync time functionality. The Netrounds Test Agent initiates TWAMP UDP streams towards TWAMP-capable routers or other devices, which reflect the streams back to the Netrounds Test Agent.

> *High throughput limits: Software vs. hardware loopback*
>
> All test methods described so far, ICMP Ping, UDP Echo and TWAMP, share one limitation: the implementation of packet loopback in routers and other reflector devices is not optimized for high data rates. This is because the loopback involves the device CPU, or in other words, it is handled in software, something which limits the attainable throughput.
>
> In activation testing, on the other hand, it is often desirable to achieve very high data throughput. TWAMP would seem to be a natural choice for this kind of test, but in practice it is often incapable of reaching the desired data rates because of the software packet loopback limitation described.

UDP Loopback: A New Measurement for High Throughput

In order to take data rates into the hundreds of Mbit/s, or even higher, it is imperative to have packets reflected in hardware. With such a procedure, however, the reflected packets will not be TWAMP compliant. Netrounds has therefore designed a new *UDP loopback* measurement tool, where UDP packets are pushed from a Test Agent towards a reflector device. The reflector device loops each UDP packet back to the Test Agent in hardware, thereby enabling data speeds on the order of 1 Gbit/s. Special configuration of the reflector device is required for hardware loopback.

Because of the need for special configuration of the router, UDP loopback is most suited for time-limited activation tests and troubleshooting. For continuous service quality monitoring, a different measurement which does not require any special reflector configuration, such as TWAMP, is more suitable.

---

[1] https://tools.ietf.org/html/rfc5357

Note: UDP loopback allows you to test up to line-rate for the majority of Cisco routers, despite the fact that standard built-in reflector features have a software limitation. Please contact Netrounds for further information on configuration.

QoS Testing

With UDP loopback and TWAMP, you can set an "Expected DSCP" parameter to test whether QoS is being handled correctly, that is, whether traffic is assigned to the correct quality-of-service class and given the appropriate priority. If the received DSCP value does not agree with the expected value, an error is indicated in Netrounds. This validation of DSCP values is important both at instantiation of a service and throughout its lifetime. A prerequisite for successful QoS testing is that the Netrounds test traffic must pass through the network elements that control QoS class assignment.

Active Measurements versus Passive Probing

Similar to active reflector-based measurements, passive probing (or monitoring) does not require the installation of hardware or software at the customer location. However, to accurately measure quality of service from the customer perspective, passive probing is a non-optimal solution for the following reasons:

1. It cannot be used for activation testing, where active traffic generation is required to test configured services in production environments prior to customer hand-off.
2. It does not monitor the service *proactively*, and can only be used for after-the-fact monitoring. Passive monitoring requires the customer to use the service so that traffic logs and counters can be collected. This will result in the customer experiencing any service degradations or problems first-hand, affecting overall customer satisfaction and requiring resources from your customer service desk for any complaints.

Summary

Netrounds' reflector-based measurements allow for the quick, easy and cost-efficient implementation of large-scale tests without the need for additional investment and configuration.

Reflection-based measurements may be carried out for service activation testing, ongoing monitoring, and remote troubleshooting, providing you with a range of benefits. Further time and cost savings can be realized when automating reflection-based testing, monitoring and troubleshooting through Netrounds' programmable APIs.

Utilizing reflection technologies for these measurement activities will allow you to:

- Save costs associated with hardware investments, field testing and technicians, and customer service desks and the associated customer dissatisfaction that comes with service quality problems.
- Decrease time to deployment of new, assured services and decrease time to issue isolation.
- Improve customer satisfaction through proactive problem resolution, increased visibility into customer problems and significantly decreased ticket handling times.
- Future-proof your network – deploy a solution that fits your physical, hybrid or virtual network needs today *and* matches your current or future requirements for automation and service healing.