White Paper: Troubleshooting Remote Site Networks – Best Practices

Management and remote site employees expect the same level of network service as the headquarters site. However, when IT staff are faced with limited resources to support remote site networks, often the applications, services and performance at those sites are not as robust as the headquarters site.

See how to deliver a high level of network service at remote sites using best practices outlined in this white paper.

TABLE OF CONTENTS

» Introduction
» Best practices for troubleshooting
» Network discovery
» Baselining
» Proactive tasks
» Reactive tasks
» Maintenance window tasks
» Solution: Integrated analyzer

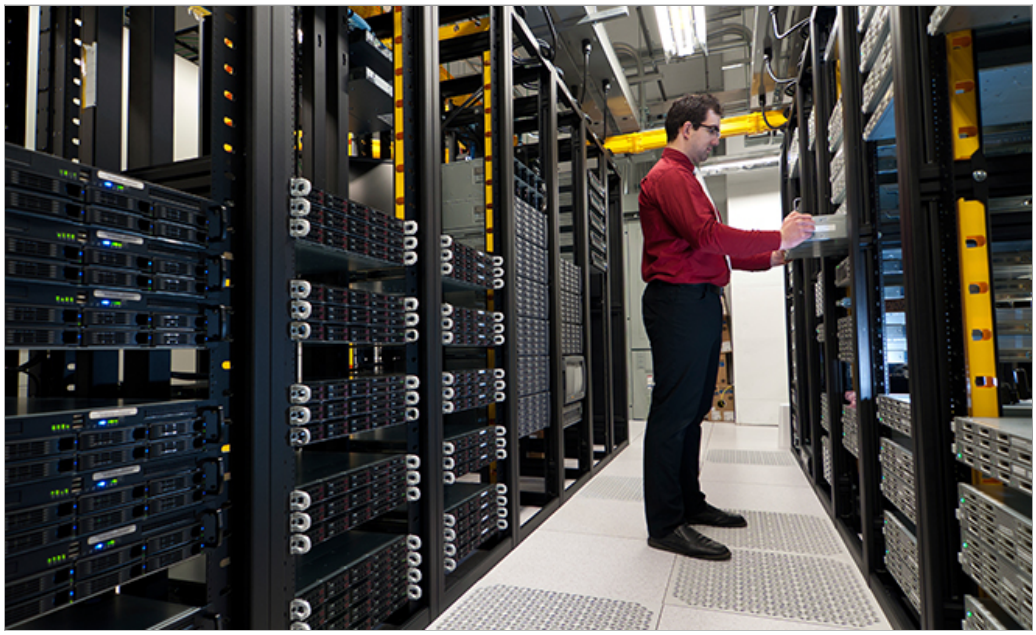

Introduction

Branch offices are a fact of life for most organizations, and while remote employees expect the same network applications, services and performance as employees located at corporate headquarters, most IT organizations do not have the required budgets or headcount to staff those sites locally. In many cases, the size of the site does not justify on-site staff, but the employees there require the same level of service as larger sites.

Although server virtualization, consolidation and the move towards web-delivered applications have business benefits, optimal productivity can still only be achieved when the same level of services are available across the enterprise. Unfortunately, even the best-planned deployment may leave remote offices and users vulnerable to performance degradation and availability issues. This creates additional challenges for the headquarters IT staff in maintaining remote site performance, availability, security and visibility.

In the headquarters environment, when remote users complain about poor performance or VoIP quality, IT staff must be able to determine the root cause of the problem and correct the situation – fast. Remote office network outages and slowdowns are made far more difficult to solve because of the challenges presented by distance, travel time and the need for tools that may not necessarily be available at the remote location. Organizing the necessary tools and dispatching staff to remote locations for troubleshooting is both time consuming and expensive, and the time spent away contributes to delays or stoppages in other critical work, such as planning and implementing IT projects.

What is needed? A strategy, process, and toolset that spans both the remote site and the corporate headquarters site. With the right information and tools, your IT staff is able to understand and resolve issues quickly and efficiently. Adding the appropriate level of visibility, IT staff could even identify remote network degradations before they become significant problems at that site. This strategy provides IT staff with the opportunity to take proactive action to eliminate congestion, latency and other problems that could affect remote users and interfere with operations. Additionally, the ability to enable staff to resolve problems from the headquarters site will avoid the need to dispatch staff, resulting in time and travel expense savings, increased network availability, and more time for mission-critical projects.

Best practices for remote site troubleshooting

Baselining, Assessment and Documentation

To gain efficiency later on, a proactive first step must be taken to establish a baseline of the existing remote site network, so that IT staff know what they are dealing with. This is especially relevant when organizations have made acquisitions or mergers where the remote site equipment configuration and network design is disparate from the headquarters equipment and network design.

The first task is a comprehensive discovery and documentation of the remote site network. This entails not only what kind of equipment exists, but also identifies the users and how and where are they connected to the network. Discovery must include information on hardware inventory, servers, access points, switch and router configurations and network connection paths. Updated maps are an essential element of “knowing” the remote site and for reference when future problems emerge.

The next necessary step is to understand what “normal” traffic levels are at the remote site. This provides a reference to work from when determining abnormal activity and compare against when validating problems in the future. IT staff must evaluate current network performance, including traffic patterns with protocol and application usage, bandwidth utilization, Internet/WAN connectivity performance, and last but not least, potential network vulnerabilities.

A good practice is to monitor traffic in and out of the site for a sample period of normal business activity. Mirroring traffic (or SPAN) from the ingress/egress router port to an analyzer capable of line-rate traffic analysis is one methodology. Another method, which also enables future visibility, is the installation of a network TAP so that engineers have a fast point of access to network traffic in the future.

Today, inter-switch trunks are widely deployed and access trunks to the desktop are very common, especially in VoIP deployments which support multiple broadcast domains together with both untagged and tagged traffic. It is therefore necessary to be able to detect all VLANs on a link and measure the traffic distribution across all those VLANs. In addition, capturing or documenting traffic statistics on a specific VLAN to allow discovery, generate traffic and capture traffic only on that selected VLAN is essential to identify protocols, top hosts and conversations limited to that particular VLAN.

Just as important, assessing Internet/WAN connectivity and quality of service(QoS) provider links provides the last piece of the essential remote site baseline. Conducting an active performance test (where test traffic is generated from the headquarters site to and from the remote site) reveals the levels of packet loss, latency, and jitter, and assurance of working QoS configurations end-to-end. Saved reports of test results can become an essential known metric from which future degradations can be compared.

HOW TO:
To provide the required information it is necessary to use many different functions and, without specially designed tools, many different products. To summarize, the following capabilities are required:

- *Network discovery – essential to find all devices and paths*
- *Mapping/documentation capabilities – automation saves labor and avoids error*
- *SNMP polling to baseline switch and router performance – granularity of data is essential*
- *Wire speed, hardware packet capture and protocol analysis to measure application response times without dropped packets –WireShark on laptop PCs is not suitable due to capture speed limitations*
- *Traffic monitoring to determine which protocols are used on the network – identifies who is using the bandwidth and for what applications*
- *Traffic generation and performance measurement – standards-based testing is preferable than ad hoc methods, or sub wire speed like that provided with freeware tools such as iPerf*

Next steps

Network professionals responsible for remote sites have to consider multiple tasks in order to support that site.

These can generally be divided into the following:

- Proactive tasks
- Reactive tasks
- Maintenance window tasks

Proactive Tasks

Once up-to-date network configuration diagrams are available and traffic levels and performance have been baselined, it will be necessary to automatically alert headquarters staff when overall traffic levels or individual critical switch port traffic has exceeded what is considered to be 'normal' levels. Many management tools (Network Management Systems or NMS) are capable of monitoring individual switch ports and WAN interface traffic and provide a method to determine when specific traffic thresholds have been exceeded on those interfaces, either by error rates or utilization rates. This will alert the IT staff to potential network degradations before they become significant problems at that remote site. However, due to their primary purpose of providing long term monitoring and trending, most management systems take samples that are too coarse for effective troubleshooting. When trying to determine the presence of intermittent or "spikey" bandwidth-hogging events, an analyzer with granular sampling rates is essential for problem detection and isolation. Additionally, seemingly minor problems such as incorrect subnet masks, duplicate IP addresses etc. should also be reported.

It is also necessary to monitor the protocols in use, which is especially important for the traffic traversing the WAN link. Are users consuming valuable Internet/WAN bandwidth for non-business related applications? Flow-based data (NetFlow, sFlow, jFlow, ipFIX) can be used to monitor the usage of bandwidth by application and user. Keep in mind, however, that flow-based data is another measure of USAGE; it contains no data regarding PERFORMANCE – the speed of transactions.

Unauthorized, unprotected rogue wireless access points – how are these discovered if there are no IT staff at the remote site to be able to walk around the site with a wireless network analyzer to find them? Monitoring sensors can be deployed to watch for rogue devices as well as monitor performance of the remote site WLAN, but again, this is where in-depth discovery from the wired side of the network becomes important; not only does the NMS need to discover IP addresses but it also needs to discover the associated MAC addresses and decode the manufacturers' prefix. Then by sorting the discovery database by MAC address, it is easy to scan the list and look for MAC prefixes that are not normally part of the network – if a suspicious MAC prefix is discovered, IT staff needs to know where that device is connected to the network (switch interface) so they can shut down the port remotely.

Reactive Tasks

When remote users complain of a "slow network", the IT staff must follow a consistent process and have access to the necessary data to identify the problem domain to identify and prove who or what is at fault. The IT staff need to quickly identify the most likely problem domain – is it the network, application, server, or client – and subsequently pass on problems with confidence by providing enough data to avoid finger-pointing and to confidently direct the problem to the responsible IT organization, but not necessarily solve an application issue.

First step – Testing Connectivity and Response Times (and the Problem with Ping)

For most network professionals, the first step in troubleshooting a problem is to ping the remote site – either the machine of the complaining user, a local server or other reliably "on" device, providing that ICMP (the layer 3 protocol used by ping) is not blocked. If ping has worked in the past, but is not now, then an examination of port status along the path is required. In the absence of "down" ports or links, an unsuccessful ping means troubleshooting from the bottom of the stack and moving up. Unfortunately, physical connectivity issues may require staff to travel to the site for troubleshooting – but do not rely on ping alone for making that determination; it may be blocked.

A successful ping at least assures physical connectivity, and can give an inexact estimation of network round trip time. But ping is not a reliable measurement method for determining packet loss, and being symmetrical in nature, provides no insight into asymmetrical link problems. Also, no user application utilizes ICMP – so whether the protocols used by a particular application can traverse the network, and the speed at which they do so, must be measured a different way – such as "opening a port". Initiating the SYN/SYN-ACK/ACK "three way handshake" of a TCP port provides a far more reliable test of layer 3 connectivity. Even better than a port connect (which validates network connectivity and network response times) conducting and measuring an application transaction provides a more reliable method of application connectivity and response times. Certain tools can target a local or remote web server, and execute and measure an HTTP GET command as a way of measuring performance of a web-based application, for example.

Keep in mind that conducting these tests from the HQ site to the remote site may provide different results than those experienced FROM the remote site.

Second Step – Examining Network Usage

It is very common for performance slow-downs to be caused by over-utilization of network bandwidth. While most LAN connections exceed the available WAN or Internet bandwidth by some multiple, it is not impossible for a local LAN connection to become overloaded, particularly if configurations are not achieving maximum throughput. Many a network engineer has been surprised to find 10Mbps half-duplex links in operation where 100Mbps full duplex or Gig was expected. SNMP or flow-based data can be examined to determine interface utilization. Granular measurement can indicate when spikes of usage are occurring, with flow data providing evidence of who is doing what.

Third Step – Testing Network Quality

A unique methodology for testing the available bandwidth is to conduct a performance test from the HQ to the remote site. Software agents are available which can be deployed on remote PCs and then targeted by an analyzer at the HQ. "Layering on" a stream of test traffic to/from the remote site provides instant insight into the quality of packet transmission, revealing issues with latency, loss and jitter that could be impacting application performance.

Fourth Step – Packet Analysis

Still operating from the HQ, the network engineer can place their analyzer in-line with the traffic feed from the remote site (either using an analyzer capable of in-line analysis, or via a SPAN port or a network tap). Keep in mind that hardware-based tools are essential for zero packet-loss analysis. The worst waste of an engineer's time would be to capture only parts of the traffic to/from the remote site, and (at best) be left guessing or, even worse, to mistakenly troubleshoot "lost packets" when the loss was actually CAUSED by the analyzer itself!

With capture files of traffic to/from the remote site, the engineer can examine delta times between frames and distinguish between network transfer time and client response time, thereby validating whether there really is an issue with performance to the remote site, or whether the issue is with the client or the HQ side.

Analyzer at the remote site

Despite these best efforts, and as pointed out already, testing from the HQ can only go so far, and is only providing test information from the point of view of the HQ site. At some point, measurements must be taken from the remote site – from the point of view of the affected users. While remote desktop (RDP) can be used to take control of a remote PC and conduct various command line tests (such as ping or tracert) these have their limitations as already discussed. The ideal scenario is to have a dedicated network analyzer on site for local testing (or, ship said analyzer to the site) but control that analyzer remotely from the HQ – eliminating the need for travel to the site.

Key to remote user application performance, and to assist in identifying the problem domain, a network services test must be provided from the remote site to ensure that vital network services are available and operating correctly. These services at a minimum would be DHCP, DNS and 802.1x authentication (if used). The effect of DNS on the performance of applications cannot be overstated. Rather than simply RDP into a client PC and conducting a CLI DNS lookup, the capability of simultaneously testing multiple DNS server addresses is necessary to perform both address-to-name or name-to-address resolution tests especially when applications are hosted on multiple servers at the headquarters site that use “round robin” DNS services for load balancing and look-ups.

Once basic services and application connectivity are validated, the analyzer must be capable of providing in-depth analysis at the remote site in order to identify the root cause of the problem. Some problems encountered at remote sites can also be intermittent and recreating those problems is getting more complex and in some cases may be impossible – if you cannot reproduce the problem, would it be safe to say that no problem exists? Unfortunately, not – it is often difficult to determine what happens on the wire, at line rate, when application error messages are received. So, there is a need to provide a capability to capture traffic that is more relevant and analyze the data when time is available, not necessarily when problems occur. In order to solve these problems and to speed troubleshooting, triggers that stop or start capturing when an event is detected both save time and provide more flexibility through:

- Unattended monitoring – capture the traffic whenever the event occurs
- Minimizing number of captures required by ensuring the event is captured the first time and avoid doing random traffic captures that may not contain anything of interest
- Analyzing the captured traffic when time is available, not necessarily when the event occurred
- Capturing traffic before, after or around the event, and only as much as needed by using capture filters to limit the amount of traffic captured and avoid having to review megabytes of traffic data.

Maintenance Window

During network maintenance times, ensure that the Internet/WAN links to the remote sites are capable of supporting the allocated bandwidth and providing quality transmission of application traffic. In order to perform this task, a network performance test (NPT) should be run between an analyzer at the remote site and a similar analyzer at the headquarters site. The test needs to be performed at various traffic rates and different frame sizes to determine if the WAN link is capable of handling the traffic, to determine packet loss and more importantly, in which direction the packets are being lost. If there are dropped packets, or the link will not support the advertised data rate, the analyzer needs to have features available to diagnose the source of the problem.

But testing for throughput and loss is only one dimension of network quality. Latency and jitter must be measured, and jitter must be measured asymmetrically if one is to understand its impact on streaming applications and VoIP quality. Also, QoS must be tested by passing traffic at various QoS settings to ensure proper traffic prioritization and prevent improper discarding or throttling of application traffic.

Where to start

As with any best practices, they are only effective if implemented and practiced routinely. The first step is to assess your current processes, priorities and needs. In addition, it’s important to understand what tools are well suited to successfully achieve the benefits of adopting these best practices. While there are tools in the marketplace that can help with some of these practices, there is just one tool that delivers full functional support for all of these best practices. The OptiView® XG Network Analysis Tablet from Fluke Networks is designed to address even the most complex troubleshooting challenges in today’s distributed network environments - fast and efficiently.

How the OptiView XG Network Analysis Tablet makes managing remote sites easier

All the functionality of multiple tools is combined into one device, making remote site management and troubleshooting easier and faster when engineers no longer have to switch from tool to tool to conduct a full array of tests. In addition, network professionals can conduct all the necessary tests at the remote site without ever leaving the headquarters site. Just plug the OptiView analyzer into your network at the remote site and you'll get 24/7 visibility into the network – it's like having a "virtual network engineer" on site. And more than one person can view the data; with the OptiView analyzer, network professionals can also work together when some staff members are off-site because data can be shared by launching multiple user interfaces simultaneously for assisted analysis and collaboration during implementation.

The OptiView analyzer then provides discovery information on network and device problems and identifies protocols in seconds. It also speeds reporting for complete infrastructure documentation. With the OptiView analyzer, network professionals can conduct a complete inventory of all network devices, where they're connected, and which services are running on them. It can do automated mapping, creating maps of the network in its current state and export that data to Microsoft® Visio, so network professionals get the data in a familiar format which can easily be used when troubleshooting the remote site. Using the OptiView analyzer, network professionals can verify and prove network readiness for network expansions, mergers, consolidations, and upgrades. They can validate and document performance, and verify new configurations to ensure the stability of the network. And they can use the OptiView analyzer to identify VLAN configurations, validate network health, audit switch/router configurations and performance.

The business benefit

Effective troubleshooting not only reduces time and travel expense but done properly, can help avoid or reduce additional hardware expenses, purchasing excess WAN capacity, unnecessary investment in outsourced troubleshooting, or having persistent problems that suck time and money from IT and the entire business.

Contact Fluke Networks
Phone: 800-283-5853 (US/Canada) or +1-425-446-4519 (other locations)
Email: info@flukenetworks.com

The business case for an integrated network analyzer

The OptiView XG Network Analysis Tablet helps network professionals manage IT projects, solve network problems and support IT initiatives, resulting in reduced IT costs and improved user satisfaction. It gives you a clear view of your entire enterprise – providing visibility into every piece of hardware, every application, and every connection on your network. No other tool offers this much vision and all-in-one capability to help you:

- *Deploy new technologies and applications*
- *Manage and validate infrastructure changes*
- *Solve network and application performance issues*
- *Secure the network from internal threats*

It shows you where your network stands today and helps you accurately assess its readiness for the changes you need to make now and in the future. Leverage the power of OptiView to give you vision and control of your network. To learn more, visit www.flukenetworks.com/optiview