NETSCOUT®

| WHITE PAPER |

OPTIMISING BANDWIDTH: WHY SIZE DOESN'T ALWAYS MATTER

ENTERPRISE

OPTIMISING BANDWIDTH: WHY SIZE DOESN'T ALWAYS MATTER

Today's network managers have to balance two conflicting demands: the desire for increased speed of response and optimal end user performance, and the need to reduce operational costs.

Traffic – business or otherwise – tends to expand to fill the available bandwidth, degrading performance across the most congested links. Network managers cannot simply throw bandwidth at the problem, but have to demonstrate that a link is experiencing congestion for a significant amount of time due to legitimate business usage before they can add additional capacity.

Approaches such as average or peak utilisation and traffic totals provide limited insight into congestion. A more effective approach is to take a burst-style view to show which links spend the most amount of time in a congested state and for how long. The network manager can then investigate why the link is busy and whether additional bandwidth is needed, as well as identifying any links where capacity could be reduced. They may be able to improve the end user experience by using existing bandwidth more effectively.

Why capacity planning is becoming critical

Capacity planning is not a new science. Network managers have always had to balance the conflicting demands of providing a robust, reliable and responsive service to their users with the constraints of their budgets. However, this challenge has now taken on increased importance as the network has become a strategic asset, supporting business-critical applications, providing the data on which decisions are made and enabling remote workers to carry out their roles effectively.

Increased use of bandwidth hungry applications such as video and Voice over IP (VoIP), the spread of virtualisation and the growth in wireless use are all making competing demands on capacity, while mobile users need to access corporate applications and resources from any location. Traffic always tends to expand to fill the bandwidth available, whether or not it is crucial to the business. Use of business critical applications is difficult to measure, and some applications may be used for both recreational and business use. For example, YouTube is an excellent medium for training videos, but is also full of amateur footage of cats, dogs and cute children!

BYOD adds to the problem by introducing devices which are outside the control of the IT team, and which may use capacity for applications and activities not related to the business. Organisations will of course address this to some extent by their IT policy, which should state what behaviour is not appropriate, but they also need the right tools in order to differentiate between different types of traffic and identify and address non-business use when it impacts on productivity.

At the same time the economic climate mean that budgets have been reduced and there is much tighter control on operating costs. Capacity is no longer a commodity, and managers cannot simply throw bandwidth at problems. Every increase has to be justified with a business case, even upgrading a single link in a large enterprise network, and making that business case takes a significant amount of time and resources. In some areas of the world, such as the Middle East, there is continued upward pressure on price, and new circuits can take anywhere from 30 to 90 days to deploy once an order is placed.

As a result, the pressures on network managers are increasing exponentially. Users expect rapid response times and put constant pressure on IT to increase bandwidth to provide an effective service to the business, while the IT department may have to report application performance and meet SLAs set by the business. They have to deliver an optimal user experience while controlling or even reducing operating costs. So it is important to be absolutely sure that any increase in bandwidth is genuinely needed and can be deployed in the right place.

Five principles for effective capacity planning

Capacity planning is not as straightforward as simply identifying which links are congested. Network managers need to understand traffic flows across the network and optimise the use of each link to provide the best possible performance with the resources they have available. They need detailed information so that they know the right time to increase throughput and can make an effective business case for change. In some instances, they will not need to increase capacity: it will be possible to improve the end-user experience and reduce operational costs by simply using existing bandwidth more efficiently.

To optimise use of existing capacity and make the business case for more bandwidth, network managers need to ask themselves a number of questions, including:

- Which network links are experiencing congestion for a significant period of time?
- Which links are most critical?
- Is the congestion caused by legitimate business usage, not recreational or 'rogue' traffic?
- Can a link be downsized while maintaining business service quality?
- Can I demonstrate that an increase in capacity in a specific link is warranted?

Before looking in detail at how to measure congestion, we will discuss five key principles to consider when tackling capacity planning.

1. Understand bandwidth resources and performance trade-offs

Bad performance does not necessarily mean that the bandwidth is not sufficient. Knowing how busy links are and for how long is key to gauging the correlation between bandwidth and performance. Performance of business critical applications may be within expected parameters for 90 per cent of the time, but no business wants to experience degraded performance for the remaining 10 per cent. Keep in mind that network bursting is normal; it just needs to be within proper thresholds. Under-utilised links can also drain bandwidth resources by using up valuable budget that could be allocated to other, over-utilised links, so identifying these links is an important part of capacity planning.

2. Use the right tools and techniques to measure utilisation

Trying to detect over-utilisation of bandwidth can be difficult when the tools are not well-suited to the job. It can be costly, timeconsuming and sometimes (depending on the tools used) almost impossible to represent congestion correctly over long time periods.

Viewing long term usage flattens out peaks of high utilisation, hiding the real problems. Peak utilisation views show which links are busiest, but do not indicate for how long. Traffic totals per day, per month etc. can show general growth in use but ignore the difference between different times of day. The key questions are: has the link been over-utilised, for how long, by what application and by which end user?

3. Account for business hours

While a network link may be busy during the night or weekend while back-ups or other updates are carried out, its performance may be acceptable during the business day. This will of course depend on the business – we work in an increasingly globalised world, and organisations also need to consider processes such as uploading retail information from stores to a central hub at the end of the working day, which may not be business critical but may impact on customer service if they are not completed on time (e.g. stock levels may not be updated).

However, it is important not to let evening and overnight data cloud your view of utilisation. By having a combination of real-time and back in time data views, IT can see what is happening across the network at different times of day, solve problems quickly and move on to more strategic initiatives.

4. Distinguish and discourage social use of bandwidth

There are two types of traffic: business and recreational. Obviously, business traffic has priority, so it is important to know why a link is busy. Is it use of a business application, or the breaking news story everyone is streaming to their desktop? Social use needs to be identified and discouraged to ensure that capacity is available for business use, which means having the right tools to identify which applications and users are causing the congestion.

Even if congestion results from use of a business application, we need to have a detailed view of how bandwidth is being consumed, when and by what. Or is the bandwidth being used by unauthorised or recreational applications that IT needs to remove from the network? Efficient application design and WAN optimisation are examples of strategic decisions that should be considered alongside the tactical approach of bandwidth needs.

5. Streamline the analysis to optimise bandwidth use and encourage forward planning

With networks growing quickly, the job of understanding which links are busy, when and why becomes more complex and timeconsuming.

The main reason that proactive capacity planning is not done is the amount of time it takes. The network team should not waste time looking at links that do not require attention, but focus on the critical few links that are busiest for the most amount of time. By implementing customised alerts that can show when bandwidth hits a pre-determined baseline, they can make the planning process easier and faster.

The limitations of traditional approaches

As mentioned in the previous section, it is important to understand how the various tools available measure congestion, over what time period, and whether they can separate inbound and outbound traffic. There are three common approaches to managing network capacity.

All of them have their limitations.

- Traffic usage totals: this approach makes it easy to display all links on a single view, highlight the ones with the most traffic and even select periodic trends to show month by month usage. However, it does not give any indication of congestion except in extreme cases.
- Average utilisation: this approach is more informative, as it will show the long term trend of utilisation, but spikes of high utilisation and even brief periods of congestion are hidden by the highly aggregated averages.
- Peak utilisation: this shows which days in the month had at least one busy minute, but does not provide any insight into the amount of time during which a link is congested. For example, it may show that peak use was on a Tuesday afternoon, but was this for all afternoon or just a few minutes?

Truly understanding network utilisation requires a high degree of granularity, but obtaining this fine detail over a month or a year requires a vast amount of data to be stored and displayed.

The solution: measuring burst utilisation

A more effective approach is to use a technique called burst utilisation. Instead of using data percentiles, the system congestion levels for each link can be reported based on 1-minute granularity regardless of the reported timeframe (day, month, or year) without loss of information fidelity. Combined with a visual display which shows >80% use in red, the network manager obtains a clear visual depiction of which links are experiencing congestion at any given time.

The advantage of using burst utilisation to measure performance is that the network manager does not have to make any decision about which data to keep and which to discard. Instead of having to work with averaged out information, or discard some of the detail, all the data is retained, and the network team can simply drill down to obtain the information required.

The network manager can then easily determine how long the congestion of a link exceeds any chosen threshold. There will always be times when capacity on a particular link reaches 100% or above. However, this may just be a slight spike in use, so will not have a significant impact on performance.

Today, with capacity on demand, it may be possible to bring a virtual server online to support occasional burst traffic. If the network carries VoIP, network managers can also implement different classes of service. Most networks use QoS technologies to protect and prioritize their VoIP traffic by tagging traffic at the device level with a queue marker (a Differentiated Service Code Point or DSCP) and setting parameters for how devices in their network treat such traffic. It is usually allowed top priority in being forwarded through the device, as well as some type of rate limit to ensure that data applications continue to perform at the levels users expect and demand.

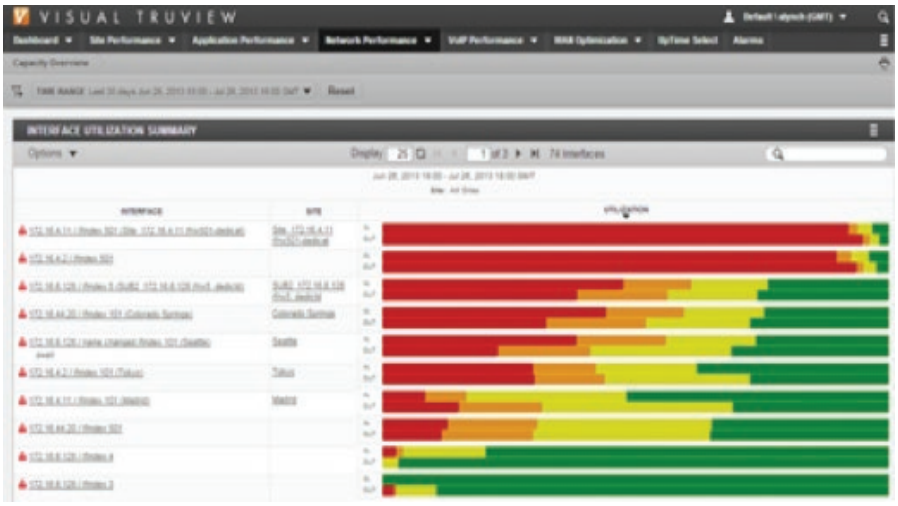

Figure 1: links >80% utilised can be clearly identified.

The team will typically want to begin to keep an eye on a particular interface when it spends more than 10 percent of time above 80 percent utilisation - which translates to a little more than a half day out of a typical workweek. When the burst reaches 20 percent of time spent at the 80 percent threshold i.e. a full working day, then it may be time to investigate how the link is being used before considering whether it needs to be upgraded. The display shows both inbound and outbound traffic.

The burst utilisation ordering can also be reversed to show which links never exceeded 30 per cent utilisation. For these links, the network manager can review whether to downgrade the link to reduce costs without any impact on performance. The money saved could then be used to upgrade the most congested links.

Ensuring usage is business critical

TruView™ is NETSCOUTs' flagship AANPM product and represents the easiest-to-use, most powerful application and network performance management solution on the market. TruView presents a unified, time-correlated view of all aspects of application component performance, allowing troubleshooters to easily find the cause and effect of performance issues with just three clicks of the mouse.

The granularity of burst utilisation information means that it can be used to examine daytime and off-peak and/or evening traffic levels.

By eliminating distortions in congestion levels resulting from nightly backup and batch jobs, the tool will show whether a link is congested during peak working hours and hence is likely to be slowing down business critical activities.

The experienced network manager will know his or her network and will understand that if there is always a problem on the final Thursday of the month, this is likely to be due to completion of monthly accounts. They can then consider what can be done to address the issue without necessarily increasing bandwidth.

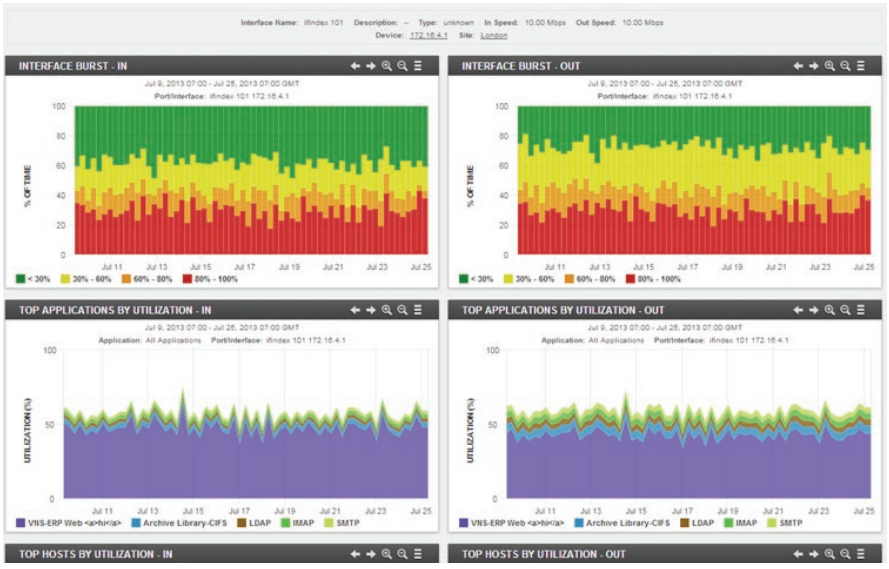

Figure 2: View utilisation levels over time and investigate who and what is making links busy.

If the congestion is only taking place overnight, they can investigate the type of traffic and if it is causing problems elsewhere e.g. back-ups are not being completed before the start of the next working day. This may indicate that, rather than adding capacity, the IT department needs to review its back-up and storage needs and introduce techniques such as deduplication.

The burst detail can also be broken out as a time trend to show if usage is stable, erratic or growing. For links which are not currently overutilised, the network manager can use >60% levels to help build long-range capacity growth forecasts and hence develop budget estimates.

Network managers can also use burst utilisation to find out if congestion is caused by valid business usage or unauthorised or recreational use.

This is done by combining SNMP data, which shows how busy a link is, and NetFlow data, which provides telemetry from routers and metadata which has a record of every 'conversation' across a link. It enables the network manager to see not only when a link is busy and how busy it is, but to identify what applications are being used and which hosts, and even which user. Are these business or recreational applications? What is causing the spikes?

One of the things we do at NETSCOUT is to educate our customers on the implications of the social consumption of bandwidth and provide them with techniques that give them insight into the problem. It can be difficult to block applications such as YouTube, as they have legitimate business benefits, but a burst utilisation tool can be used to find out who is transferring 1MB of streaming data and the team can then investigate the purpose of the download.

Saving network managers time

Capacity planning can be time-consuming without the right tools in place. One of the reasons NETSCOUT added capacity planning capability to its TruView product was in response to customer needs. They told us that finding out which links were congested and needed upgrading required a significant amount of effort, and traditional products did not provide the granularity they needed. They could not see how busy a specific link was, how long it was busy for or take out data for out of hours use. For this reason we took the burst utilisation approach, as it provided the visibility, granularity and level of detail we were being asked for.

Customers tell us that they use TruView for capacity planning in two different and complementary ways. First, it is used reactively to identify problems and resolve them quickly. As soon as a user reports a congestion issue, the network team can identify the link which is the source of the problem and look to see what applications are being used, whether they are business critical and whether the problem is transient or longer term.

Secondly, it can be used to proactively address congestion issues. The user can go to the capacity overview page, look at the top five busiest links, and then build intelligent QoS policies that make the best use of available resources, and prioritize the delivery of critical applications Capacity planning and bandwidth management may not be new, but they are more important today than ever before as the network becomes increasingly business-critical. By measuring burst utilisation and analysing the results, network managers can quickly get to the heart of the problem, ensure they have capacity in the right places and free up time for other activities.

About NETSCOUT

NETSCOUT SYSTEMS, INC. (NASDAQ: NTCT) is a market leader in real-time service assurance and cybersecurity solutions for today's most demanding service provider, enterprise and government networks. NETSCOUT's Adaptive Service Intelligence (ASI) technology continuously monitors the service delivery environment to identify performance issues and provides insight into network-based security threats, helping teams to quickly resolve issues that can cause business disruptions or impact user experience. NETSCOUT delivers unmatched service visibility and protects the digital infrastructure that supports our connected world.

Solutions from NETSCOUT

TruView

TruView™ provides the ability to track, baseline, trend and monitor individual application performance of every end user experience, enterprise-wide through a highly customizable dashboard. It also provides high volume packet archival at 10Gbps line rate and comprehensive VoIP/Video monitoring and troubleshooting.

More information at enterprise.netscout.com/truview

OptiView XG – Automated network and application analysis

The OptiView XG® is the first tablet specifically designed for the Network Engineer. It automates root-cause analysis of network and application problems allowing the user to spend less time on troubleshooting and more time on other initiatives. It is designed to support deployment of new technologies, including unified communications, virtualization, wireless and 10 Gbps Ethernet. The result is that new initiatives get up and running faster and network stay productive even in these days of smaller teams.

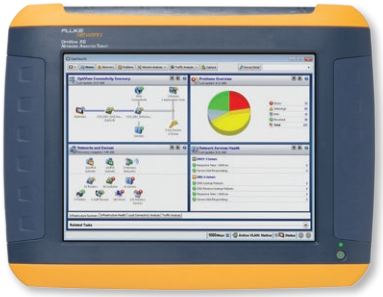

More information at enterprise.netscout.com/xg

OneTouch AT – An automated all-in-one tester for understanding end-user network performance on-site in real-time and through the cloud over time.

The OneTouch AT™ Network Assistant is an all-in-one Gigabit Ethernet troubleshooter for copper, fiber optic and Wi-Fi networks. It provides a client view of network performance so you can resolve problems fast and complete deployment projects on time.

More information at enterprise.netscout.com/OneTouchAT

NETSCOUT

Americas East
310 Littleton Road
Westford, MA 01886-4105
Phone: 978-614-4000
Toll Free: 800-357-7666

Americas West
178 E. Tasman Drive
San Jose, CA 95134
Phone: 408-571-5000

Asia Pacific
17F/B
No. 167 Tun Hwa N. Road
Taipei 105, Taiwan
Phone: +886 2 2717 1999

Europe
One Canada Square
29th floor, Canary Wharf
London E14 5DY
United Kingdom
Phone: +44 207 712 1672

NETSCOUT offers sales, support, and services in over 32 countries.

For more information, please visit
enterprise.netscout.com

6003157_EN