



Jam today - but not tomorrow?

If you want to keep traffic moving across your network you could build extra lanes. But will buying more bandwidth really solve your congestion problems?

If your WAN's performance is getting clogged up by too much traffic, the solution is simple: just add more bandwidth, right? That question is, of course, as loaded as your network and is universally met by industry experts with a resounding (and often scathing) 'no'.

"When people find that their applications are running slowly across the network, having hopefully first eliminated their servers as the problem, they turn to their network. Often they seek advice from their ISP and, surprise surprise, they are recommended a capacity upgrade," says Frank Puranik, director of iTrinegy which specialises in tools to measure and optimise application performance in networks. "The truth is that they may not need this. And if they do get more capacity it can be quickly eaten up."

The fact of the matter is that adding bandwidth to address poor WAN performance can only help to resolve part of the problem. Network delays and data delivery problems are caused by a mix of bandwidth, latency, congestion, protocol, application and data size. Adding more bandwidth just addresses one of these issues and only during the last mile, as Nigel Hawthorn, VP of Marketing LMLA for WAN optimisation specialist Blue Coat Systems, points out: "A simple analogy would be if you need to drive from London to Leeds. In pick up a parcel and return. Would adding an extra lane in the last mile between your house and the beginning of the M1 make much difference? The distance is still 190 miles or a round-trip of 380 miles.

"Our computer protocols would probably insist that you split the parcel into 1,000 small pieces and drive backwards and forwards picking up each piece one by one. Only then would you find you

already own the great majority of what you'd already picked up and that a portion of the trip wasn't necessary."

So instead of 'building more lanes' by buying more bandwidth, network managers should look at optimising the WAN.

Continuing his analogy, Hawthorn says that such solutions can decide if a trip is needed at all, and only if a portion of the parcel is actually new. "[They] compress the parcel into a smaller total, and allow you to collect multiple parts at once making a much greater difference than just adding bandwidth for a small portion of the journey."

What are you carrying?

Rob Marson, Strategic Marketing Manager at network technology specialist JDSU, says WAN congestion has increased as enterprises reduce the number of servers and application instances from remote branch offices, consolidate data centres, and centralise applications. "Many applications are very 'chatty', exchanging multiple messages or acknowledgments between sender and receiver. Though not an issue in LAN environments, when traversing the WAN, the increased latency and the congestion caused by transaction-heavy applications results in poor throughput and slow application response times. This results in user frustration, reduced productivity, impacts business, and any intended cost savings could be negated."

Silver Peak is one of the big players in WAN optimisation and its VP of marketing Jeff Aaron says that bandwidth is only a third of the problem addressed by WAN optimisation. "There are also latency issues due to distance, and quality issues due to congestion (i.e. packet loss). Buying more capacity only helps when bandwidth is the limiting factor. If latency or loss are the limiting factors, then more capacity won't make applications run faster or with higher quality.

"Also, in many instances, more bandwidth isn't necessary what's

needed is better bandwidth utilisation. By properly prioritising traffic and allocating resources, network managers can see significantly better performance with the bandwidth they have available."

But not everyone is so dismissive of adding more capacity to alleviate WAN congestion. Talari Networks has developed Adaptive Private Networking (APN) technology for WAN virtualisation and its marketing VP Keith Morris says that the decision to add more bandwidth depends on the application. "WAN optimisation solutions help when latency is the issue or simple data can be compressed.

However, there are a growing number of applications that simply need more bandwidth. This is especially true as more applications are centralised in fewer places as part of a private or public cloud. "You really need to understand if your performance issues are related to bandwidth or just the fact that you are trying to use a protocol over the WAN that is very chatty and that was never really designed to work well when there is a large amount of latency. The classic example is Microsoft CIFS which is



"Is it time to discard WAN optimisation appliances in favour of higher capacity WAN connections?"

*Rob Marson,
Strategic Marketing manager,
JDSU*

known to perform poorly over the WAN. Then there are applications that are not necessarily impacted by latency but don't have a good user experience if there is insufficient bandwidth - applications such as VoIP, video, virtual desktops, and the transfer of already highly compressed data such as design files or images."

So when it comes to WAN optimisation the first thing that IT managers need to do is to understand what is using the network. Buying more capacity is sometimes the solution to congestion, but not when latency is the problem, according to Paul Moorhead, head of Product Engineering at Replify which provides WAN optimisation products based on software-only virtualised appliances. "It's latency, not bandwidth, that can be throttling many user activities. Another way to put this would be to say that more bandwidth will allow more users to have the same horrible experience, but it won't allow any individual user to have a good one."

Traffic management

So what does WAN optimisation actually do? Hawthorn says such technologies reduce the overall amount of traffic that is transferred, compress the data that has to be transferred, and reduce the serial nature of many protocols that induce delay by waiting for acknowledgments. He adds that WAN optimisation can be implemented in dedicated appliances, virtual appliances, and client software for devices.

Marson says that traditional WAN optimisation appliances were very popular in the early days of enterprise IT transformation as network architects first transformed their WANs to IP, and then focused their attention on reversing server sprawl.

"Typically limited by the 'skinny pipe' WAN bandwidth options offered by traditional service providers, WAN optimisation appliances played an important transition role for many IT administrators. In the last number of

WAN optimisation

register online @ www.networkingplus.co.uk to receive your free monthly copy

years, as IT strategies have evolved to consider both disaster recovery and cloud computing initiatives in combination with increased choice and flexibility in terms of the options provided by service providers, enterprise IT architects are re-thinking their WAN strategies. Is it time to discard WAN optimisation appliances in favour of higher capacity WAN connections? Or do both co-exist as part of an end-to-end strategy?"

In Marson's view, WAN optimisation suppliers tend to target organisations with remote branch offices and aim to improve application performance by reducing transactional latency over the WAN through protocol spoofing, caching, and data compression. "Typically, these appliances were implemented using software based in general purpose CPUs. Targeting skinny pipes such as T1 or T3, they lack the performance capability to achieve multi Gigabit Ethernet wire rates. Often, the reality of implementing a solution involving these types of appliances is complex as their implementation is highly specific to the application.

"For instance, if the software developer changed the characteristics of the application (i.e. adjusting the number or type of messages a transaction would use) this would require an associated upgrade to the WAN optimisation appliance. Similarly, the appliance vendor has to spend considerable effort and expense developing and qualifying their appliance with various versions of applications. Their cost, as well as the complexity to deploy and maintain, requires strong oversight by the enterprise IT administrator."

Essentially, it's all about traffic management, and when it comes to WAN optimisation, there is a general consensus that without Application Performance Management (APM) tools, businesses will not have any insight into how network capacity is being used and will not be able to optimise it to accommodate mission-critical applications.

Dominic Jones, Product director at Cable&Wireless Worldwide, says that APM technology offers an effective route to optimising a network and significantly improves reliability. By analysing network traffic in real-time and optimising essential traffic throughout the network, it helps businesses use existing capacity more intelligently.

"APM functions much like the triage ward in a hospital, prioritising urgent patients to match the available doctors.

Just buying bandwidth would be like hiring hundreds of new doctors and paying for there to be enough of them to see every person who came into the hospital, instantly on arrival. Doing nothing is the equivalent of holding back a heart patient whilst the doctor finishes administering to a cold. APM, on the other hand, balances demand against resources and prioritises to ensure the best outcome. It adds the intelligence to be able to make choices," he says.

Can WAN congestion be avoided in the first place? Silver Peak's Aaron says that as long as you have multiple applications sharing a WAN, or are using shared WAN technologies like MPLS and cloud, congestion is inevitable. This is especially true for international WANs where traffic is traversing multiple carriers' networks. There is no way to completely avoid congestion without spending a fortune on dedicated WANs, so the key is to mitigate the impact of congestion when it does arise."

Talari's Morris agrees: "It is completely possible to mitigate network congestion. Well-designed WAN visualisation solutions shape traffic in such a way that they don't inadvertently cause 'last mile congestion' at remote sites. If any network impairments such as congestion, loss or excessive jitter are detected, these solutions can re-direct traffic sub-second to network paths that currently offer the best possible transport for that type of traffic. For example, VoIP packets are not only sent out of the box first, they are placed on the network path with the least amount of loss, jitter and latency."

iTrinegy's Puranik also reckons that congestion can be avoided but only if you monitor your network. "Without doing this you're groping in the dark. When you put in a WAN for the first time, you should measure how much traffic is going to need to flow over it. If this is a cloud migration (or data centre consolidation) you can measure what the apps are using today with a good APM and factor this in. But the higher latency of the WAN may mean that apps that worked fine in your LAN won't want to play nicely in the WAN."

Application management is clearly key here. Blue Coat says that managing the WAN is about making the right decisions and selecting the right tools for your requirements. "For example, you can choose to block access to inappropriate content such as online games during work hours," says Hawthorn. "At the same time you may want to allow some aspects of social networking sites such as Facebook, which can help to improve communication and productivity within the workplace, whilst blocking out other aspects such as gaming which has both a detrimental effect on bandwidth consumption and efficiency."



"It is completely possible to mitigate network congestion."

Keith Morris,
Marketing VP,
Talari Networks

Optimise or virtualise?

According to Replify's Moorhead, when WAN optimisation products first appeared they were "expensive tin boxes" providing site-to-site acceleration. But he notes that today's users are mobile, work from many different locations, and often without any reliable connectivity at all. "What's needed is technology which seamlessly shifts between delivering shared WAN optimisation for a user when they cluster with others, and individual WAN optimisation for them when they roam."

He reckons that Replify goes a step beyond its competitors by having a product which allows users to continue to work with shared content from file shares, such as SharePoint for instance, even when the network connection is lost or is so slow as to be useless. "We take care of synchronising local and remote copies whenever connectivity is available and the user need not know, or care, whether lie's on or offline," says Moorhead.

His colleague, UK and Ireland Sales manager Tony Rowland, explains how Replify recently worked with a client who was transferring large files up to 10GB in size from the US to the UK: "With 120ms of latency, the session was constrained to 480Kbps. With our product optimising the protocol, we increased the session throughput to 45Mbps - nearly 100 times the original throughput. Simply adding bandwidth would do nothing to increase throughput over high latency links. This is when WAN optimisation can be the very best course of action as it removes the effect latency has by optimising protocols such as CIFS, MAPI, FTP and HTTP(S) to increase throughput above and beyond what adding bandwidth will do."

So if your network is using a mix of latency sensitive and high bandwidth applications, is WAN virtualisation - rather than appliance-based optimisation - the best route to go clown? In such cases, Morris says that it makes sense to use WAN optimisation to 'fix' an organisation's chatty applications and then use WAN virtualisation to add as much inexpensive bandwidth as needed for the rest of the applications. "We see many organisations that initially deploy WAN optimisation and then move to WAN virtualisation to scale bandwidth economically. When network managers look at the cost of private WAN bandwidth, which is 30 to 100 times more expensive than internet connectivity, it makes sense to use more internet if they can buy a solution that eliminates the reliability issues associated with IPsec VPNs."

Silver Peak naturally advocates the use of appliances and says that its WAN optimisation devices have become pretty robust. Aaron says: "For instance, a single Silver Peak solution has TCP acceleration techniques to mitigate latency, Network Integrity features to fix packet delivery problems (i.e. dropped and out of order packets), and deduplication for bandwidth savings. In addition, it incorporates QoS and traffic shaping to avoid doing this in a separate device and includes IPsec to eliminate the need for separate VPN concentrators.

"Such a solution also focuses on solving problems within the network, not at the application layer. Therefore,

it works on all IP traffic - today and tomorrow. This is in stark contrast to caching solutions, for example, that only work on a single application. The result is that one product can optimise an entire global network."

JDSU's Marson adds that as WAN optimisation appliances shift to software clients they may play an important role in enabling mobile computing from smart devices or tablets. He says that the increased availability and affordability of fibre-based WAN connectivity services makes them an attractive option for connecting large enterprise sites together or to third party data centres. "Moreover, the ability to transport different protocols (such as Fibre Channel, Ethernet, FCoE) simultaneously with low latency and without data loss, makes fibre-based WAN connectivity the natural choice for enterprise IT architects when implementing cloud computing, data backup, and disaster recovery strategies." When to use what strategy depends on the specific circumstance or problem which needs to be addressed. "While high bandwidth, low latency, fibre-based WAN connectivity options have proliferated, skinny pipes still exist and in many cases are the only options available to connect remote locations," says Marson. "As a result, WAN optimisation appliances will continue to play a vital role in an overarching enterprise WAN strategy." •