



Abner Germanow
Director, Enterprise Networks

The Reliability Imperative for Enterprise Networks

May 2007

As IT managers evaluate technologies for their next-generation networks, they face growing demands for higher application availability and performance as well as greater agility and flexibility in aligning IT with business priorities. Today more than ever, enterprises need network devices that provide best-in-class resiliency, scalability, and performance in all deployment configurations. Increasingly, enterprises are turning to distributed hardware and modular software solutions to achieve more predictable application performance, increased network availability, and lower operating costs.

The following questions were posed by Force10 Networks to Abner Germanow, director of IDC's Enterprise Networks research, on behalf of Force10's customers.

Q. Why is the issue of network reliability so important throughout the enterprise today?

A. Reliable networks have always been at the top of the network administrator's shopping list, but with the introduction of voice on the network over the past several years, tolerances for latency and availability previously limited to datacenters are now required throughout the enterprise. In most cases, network managers have been planning for this shift in a variety of ways, but the reality of this traffic has started to set in only recently. Voice influences latency and availability, but other traffic types such as video — where there's tremendous bandwidth demand — software-as-a-service applications, and peer-to-peer collaborative applications are changing traffic patterns across the enterprise.

The bottom line is that click-and-wait-type applications, such as Web browsing, email traffic, and business applications are being joined by real-time communications applications such as phone conversations and video streams, where latency is now measured in milliseconds.

Q. What does "reliability" mean in the context of an enterprise network?

A. There are several different aspects to reliability. To the end user, reliability means that there's an expectation that a service is available. In this context, reliability means that the user experience on the network will consistently deliver the service the user is trying to access whenever the user feels he or she needs to access the service. Today, that means the network must be up 24 x 7 with no unplanned downtime.

Reliability also means control — the control you have over how all the different applications influence the end-user experience on the network. It's how you monitor and control the application traffic flows in a route and deliver the applications to users and devices on the network in the way that's expected.

Reliability is also connected to scalability — the ability of the network to perform as it always has even when there is a crush of traffic and to deal with the surge gracefully. In this context, scale means the ability to perform consistently, regardless of what else is happening on the network, as well as the ability to easily add capacity without having to rearchitect the network design. It means easily growing the network so that you always have enough runway.

Security is also a function of reliability. Network attacks, whether malicious or caused by human error, have the ability to impair network services. Networks need to be able to shield business traffic from malicious traffic to maintain the integrity of the data flowing across the network and application delivery.

Q. When should an enterprise deploy a multivendor environment versus a single-vendor environment?

A. The issue of when to deploy or consider deploying a multivendor environment is, in some cases, dependent on a desire to ensure that you're not reliant on a particular vendor or a particular technology direction that may not be consistent with your company's views on where your network investment should go. The general advantage of single sourcing is that you can gain a level of consistency across the network in terms of devices and interfaces, which is harder to achieve when you use multiple vendors. So when you do use multiple vendors, a best practice is to not have too many. Decide on the minimum number of vendors that you can efficiently and effectively manage, and commit to those particular vendors.

Another influence in the single-vendor versus dual-vendor decision is the need to protect yourself from cascade failures that might occur if a security exploit renders one vendor's equipment unstable. For example, if there's a security or other reliability problem with your single-vendor environment, then you may be exposed to a level of risk that is unacceptable to your organization. A dual-vendor environment can help mitigate the risk of a cascade or catastrophic failure.

In addition, if you have particular functions such as security or performance needs that can't be met just by a single vendor, then a dual-vendor or multivendor strategy can help you acquire functionality that isn't available from a single vendor. Just be mindful that as you increase the number of vendors, the operational cost of managing that environment also rises.

Q. How are operating systems and network devices evolving to meet the new demands on enterprise networks?

A. Network devices have always had operating systems, and traditionally those operating systems have been fairly well hidden from the view of network managers. Network managers interacted with the management and configuration of applications that were running on the operating system or, in many cases, that were part of the operating system.

The problem with this kind of monolithic environment, where the applications and the operating system are all one piece of software, is that eventually the software becomes unwieldy and difficult to manage — especially if you want to add new features and functions rapidly without having to disrupt the service of that network device.

Traditionally, taking a network device offline, updating the software, and bringing the network back online was an acceptable practice, provided it wasn't done very often. In other words, having planned downtime was okay. Now, however, enterprises want to be able to apply security patches and other pieces of functions while the network device is in operation.

The operating systems that run network devices today are starting to look more like operating systems that run applications in the rest of IT, where there's a general-purpose operating system and then a set of applications that load on top. This architecture allows the technology vendor to cycle new functions very rapidly, and it also eases the management of updating and changing configurations as well as the management of a consistent software base across the network.

Increased network uptime and reduced downtime essentially come from the ability to change an application on the fly, which means that, in some cases, you can load a new function onto a network device without having to stop the traffic from flowing through that device. It's analogous to newer applications like Skype, which you can update without having to restart your computer afterward. Having a modular operating system that mimics the structure used in the rest of IT enables network administrators to make updates faster than they've traditionally been able to.

Q. How can enterprises better future proof their networks?

- A. When an enterprise looks at future proofing its network, it needs to focus not just on what's supported today but also on how to evolve the network to support changes in availability requirements, changes in latency requirements, and changes in scale requirements. A fourth factor might be capacity, but scale and capacity are somewhat tied at the neck.

The key thing to remember is that when you design and deploy a network in support of what's on the network today, you don't really know what you're going to have three years from now — even when the devices you're deploying today may still be in production deployment. The lesson of the past several years is that there is very likely a network-based business in your company's future. Many enterprise network administrators can count on their network evolving from a support function to a primary business function.

For example, we've been talking about voice and video in the networking world for a long time, but in 2006, the market made a fairly major turn and really began deploying voice in earnest. If you don't have voice on the network today, expect voice to come. If you don't have significant video on the network, expect video to come.

That said, it's still quite difficult to determine how, exactly, network usage will change in the future, so don't put too much effort into determining the type of change. The effort is better spent in designing and deploying a network that requires a minimal amount of change to deployed devices when the change in traffic inevitably occurs.

ABOUT THIS ANALYST

Abner Germanow is the director of IDC's Enterprise Networks services. In this role, he oversees a team of industry experts and their comprehensive research and analysis on evolving enterprise network infrastructure markets, including wireless LAN, IP telephony, LAN switching, and enterprise routing.

ABOUT THIS PUBLICATION

This publication was produced by IDC Go-to-Market Services. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Go-to-Market Services makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

COPYRIGHT AND RESTRICTIONS

Any IDC information or reference to IDC that is to be used in advertising, press releases, or promotional materials requires prior written approval from IDC. For permission requests, contact the GMS information line at 508-988-7610 or gms@idc.com. Translation and/or localization of this document requires an additional license from IDC.

For more information on IDC, visit www.idc.com. For more information on IDC GMS, visit www.idc.com/gms.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com