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

IDC EXECUTIVE BRIEF

Achieving Enterprisewide Network Reliability for Complex, Real-Time Applications

September 2007

Adapted from *Worldwide Ethernet and LAN Switch 2007–2011 Forecast* by Abner Germanow and Petr Jirovsky, IDC #206031

Sponsored by Force10 Networks

Introduction

The demand for high-performance enterprise networks continues to grow as organizations look to add applications such as voice and video to the mix. Because of the real-time nature of these next-generation applications, organizations must find ways to increase network reliability. This Executive Brief looks at the changing nature of reliability, including demands for reduced latency, and how a modular distributed hardware approach can improve network reliability while helping organizations prepare for future high-performance applications.

The Reliability Imperative

According to IDC research, the enterprise networking equipment market will continue to grow faster than other enterprise IT equipment markets. Why? The number of users and associated devices continues to increase, the inevitable combination of voice and data over the same network will require upgrades, and companies continue to expand to meet the needs of a global, mobile workforce. In addition, the demand for reliability will increase to meet the need for real-time business-critical applications and communications.

With business moving at a faster pace, high application availability and performance — along with greater agility and flexibility to align information technology with business priorities — are imperative. While reliability has always been a critical need for enterprise networks, the growing inclusion of voice applications over the network has dramatically reduced tolerances for latency and availability. In addition, due to an increasingly distributed workforce, these reduced tolerances are concerns not only of the datacenter but also of the entire enterprise. While voice and associated user expectations of service availability and quality are critical drivers, other applications, such as video, software as a service, and peer-to-peer collaboration, are increasing bandwidth needs and dependence on real-time communications.

The enterprise network is no longer a domain for click-and-wait applications, such as email, Web browsing, and traditional datacenter business applications. Real-time communications, including voice over IP (VoIP) and streaming video, are becoming the norm. As a result, latency is now measured in milliseconds and network reliability is more critical than ever.

However, reliability also has grown to mean more than just network uptime. In addition to higher application availability and performance, today's distributed organizations increasingly demand greater agility and flexibility in aligning information technology with business priorities such as scalability, resiliency, line-rate performance, the ability to make upgrades and repairs on the fly, and increased security demands. By the way, in today's cost-conscious business environment, these attributes must often be achieved or improved with flat budgets and staffing levels.

Reliability: What It Means in Today's IT Environment

As the nature of enterprise applications evolves, so does the definition of "reliability." The term now means different things to different people in an organization. To end users, reliability is simply knowing that service is available whenever needed. Today's distributed and mobile workforce is accustomed to wireline telephone service availability, and any disruption could be more than an inconvenience when conducting business in real time.

But for datacenter IT managers, reliability means much more than 24 x 7 availability. It also means control of how all the different applications running on the network influence the end-user experience. It's managing application traffic flows to users and devices to meet their needs.

In addition, reliability means scalability, the ability of the network to continue performing within expected ranges even when there is a dramatic increase in traffic. The network must perform consistently for all users with minimal latency regardless of what is happening elsewhere. Additional applications and capacity must be added easily, without having to reconfigure or rearchitect the network.

Reliability also includes security. Network attacks, malicious or not, can impair network performance. Today's multipurpose networks must shield traffic from attacks or viruses to ensure the integrity not only of the data but also of the end-user experience.

Achieving Reliability with Modularity and Integration

Hardware reliability and connectivity, in the form of network device uptime, are now a given, with network equipment suppliers providing highly reliable devices. As a result, reliability today means managing distributed applications, the number of vendors an organization works with, and how to economically meet the needs of an expanded network. In short, reliability means consistent, real-time or near-real-time application availability.

Ironically, as applications, and associated user demands, continue to grow in complexity, reliability becomes dependent upon simplifying the network. As datacenter functions continue to push outward across the enterprise, management of services to meet end-user needs increasingly relies on standardization and flexibility.

Remote branches, for example, are increasingly recognized as a strategic location to which IT can add value to the business. This further raises staff and budget issues as real-time applications and the increased demands for bandwidth, control, and accessibility mean that every branch of an organization must be part of the network.

How can organizations add equipment, such as servers and switches, in both the datacenter and the remote office in a way that not only increases control but also makes it easier to add capacity or redundancy?

One strategy is to take a distributed hardware approach to network architecture. Using modular distributed devices becomes a less expensive way to expand a network and maintain expected levels of service, particularly when these devices integrate switching and routing.

First, this modular integrated approach reduces management costs and tasks as well as increases operational efficiency. Fewer standardized devices eliminate complexity without sacrificing performance and make it easier to ensure that users' performance needs are met. Devices can be placed throughout the enterprise, creating hardware, operating system, and application consistency. Modular systems also make it easy to increase capacity as new boards, or blades, can be plugged in and used immediately. Similarly, repair, servicing, and upgrades can be done rapidly.

Further, a distributed approach using integrated devices can change the network operating system paradigm, which will dramatically add to reliability. While network devices have always had operating systems, they traditionally were hidden from the view of network managers. Control and configuration of applications, in many cases, were part of the operating system. The challenge with this approach, where application and operating system are one piece of software, is that the software becomes unwieldy and difficult to manage. In addition, adding new features and functions required taking all or part of the network offline. While this was acceptable to update "old school" network applications, provided it was not done frequently, it's not acceptable today.

With real-time demands for communications and applications, changes must be done on the fly, with minimal interruption of service. As a result, today's network operating systems are becoming more like those that run applications in the rest of IT: a general-purpose operating system with a set of applications running on top. New functionality, patches, security updates, and so forth, can be added very quickly, not only reducing or eliminating downtime but also making it easier to update and change network configurations in a changing business environment. Most important, management is simplified due to a consistent software base across the network. A reliable, distributed network requires this modularity.

Balancing Vendor Management

As organizations strive to simplify and standardize real-time networking, they are taking a hard look at the number of network equipment vendors they work with. This often requires striking a balance between gaining a level of consistency across the network in terms of devices and interfaces, which requires fewer vendors, and making sure that the company is not reliant on a single supplier or technology direction that may not be consistent with the organization's strategic direction.

Security is a critical factor here, as a multivendor network helps protect against cascade failures if a security problem is targeted at a single vendor's equipment. Many organizational security and performance needs cannot be met by a single supplier. A multivendor, or at least a dual-vendor, strategy can help an organization acquire functionality that isn't available from a single source. However, as the number of suppliers increases, so do the operational costs of managing the environment.

Single sourcing does allow for a level of device and interface consistency across the network. As an organization looks to simplify, it must strike a balance based on the minimum number of vendors that can be effectively managed.

Planning for the Future

There's an added benefit to modularity and integration across the enterprise network — future proofing. When installing or upgrading network equipment, organizations need to focus not just on today's requirements but on how to meet changing availability, latency, and scale requirements. While it's difficult to know what will be available or needed three years from now, for example, modularity will make it easier to continue migrating the network from a support function to a primary business function.

While voice and video applications have been discussed for a long time, they are not yet in widespread use across enterprises. This situation is changing, however, and if an organization does not have voice — or video — applications today, it will soon. With these applications, reliability translates into ability: Can the network easily accommodate the demand and the required service level?

Leading the way are distributed networks for high-performance computing. These bandwidth-intensive and low-latency applications demand modularity and integration — simplification — and offer remote performance typically found only in the datacenter itself. These networks are designed for systematic upgrades, particularly increases in bandwidth and associated management and control.

One dramatic example is the movement to 10 Gigabit Ethernet (10GbE). Increasing network and datacenter server density, multicore processors, and the associated applications are forcing organizations to increase bandwidth capacity. By taking an

integrated distributed approach, particularly using modular device backplanes that are designed for high-speed communications, organizations can rapidly upgrade to provide the availability and speed required for IT consolidation, multimedia, high-performance computing, VoIP, and IPTV/video on demand. With these applications, the reliability expectations, which include security overhead, are such that high performance (10GbE) is a prerequisite to meeting the quality-of-service requirements.

Conclusion

While reliability has always been a critical need for enterprise networks, the growth of high-bandwidth applications such as voice and video over the network has dramatically reduced tolerances for latency and availability. In addition, with an increasingly distributed workforce, these reduced tolerances are a challenge not only for the datacenter but also for the entire enterprise. Ironically, as applications and associated user demands continue to grow in complexity, reliability becomes dependent upon simplifying the network. One strategy is to take a modular distributed hardware approach to network architecture.

An integrated modular approach reduces costs and increases operational efficiency by eliminating complexity without sacrificing performance. Devices can be placed throughout the enterprise, creating device, operating system, and application consistency. This approach dramatically improves reliability by making management, updates, and the addition of new applications easier.

Further, a distributed modular approach to enterprise networking can prepare organizations for future reliable high-speed, low-latency applications. These networks are designed for systematic upgrades, particularly increases in bandwidth and associated management and control, making it easier to adopt reliable 10GbE, for example. This will pave the way for even more real-time, business-critical applications.