CDW PEOPLE WHO GET IT

THE NETWORK ADMINISTRATOR'S GUIDE TO ROUTERS & SWITCHES

Strategic acquisition and management practices help ensure long-term network performance.

Executive Summary

Network design is undergoing a revolution as old assumptions give way to new realities. Enterprises that cannot strategically market themselves and communicate effectively in the Internet age are finding it difficult to compete.

While change is often slow, organizations are streamlining their processes in order to become more efficient in their service delivery and adapting to customers' demands. If correctly advised, network adaption will provide a competitive edge in any industry. That is the goal of any IT project – leveraging technology to do more with less.

Table of Contents

This paper is a starting point for enterprises considering making communication upgrades to their core infrastructures. Having an understanding of the organization's requirements will drive positive change. If done properly, communication upgrades provide a value that is important to quantify in order to ensure a positive return on investment (ROI) when purchasing the technology. Otherwise why make the investment?

Enterprise Design Considerations

Networks empower people to do their jobs, communicate and collaborate. The users of a network have very little knowledge of its interworkings and how information is relayed to their end-user device. As long as it works, the organization will have satisfied users. That is why networks are built with redundancies and why critical information is backed up.

Virtualization is changing data center performance requirements by concentrating immense power in small spaces. Mobility and bring-your-own-device (BYOD) initiatives are transforming Wi-Fi from a guest service into a mission-critical service. Cloud computing and increased reliance on Internet-based resources are making network reliability and application control top priorities in all types and sizes of enterprises. As infrastructures become more complicated, the days of being reactive are no longer acceptable for the savvy organization.

The foundational elements for all enterprise communication are routers and switches. These are the key building blocks. In all enterprise communication networks, these devices are the core infrastructure that enables smooth and reliable communication. Well-architected infrastructures are equipped to take advantage of today's requirements and are poised for tomorrow's demands.

When architecting new infrastructures, it is important to understand the technical and business requirements of the enterprise before making changes. Network architects tasked with building modern networks have no choice but to question old design rules and to rethink how they assemble these critical pieces to reach the organization's goals. It is very important to understand current trends while utilizing networking best practices. This knowledge will lead to the best results and position the enterprise to adjust quickly to market changes while maintaining its competitive edge.

Prior to selecting devices to meet the organization's needs, it is important to understand the critical protocols that are currently in use for its routing and switching. Any new devices should be designed for flexibility to enable the enterprise to take advantage of new, more efficient protocols. With this knowledge, the network engineer can select the best devices for reliability, flexibility and management. These factors will help IT professionals make the most of what they have and, when needed, build for the future.

Monitoring and Maintaining

It is also important to continuously monitor and maintain all network devices. This step is often skipped, but experienced engineers will know that a lack of ongoing maintenance can lead to catastrophic network outages.

Consultants will often categorize an organization without monitoring capabilities as a "reactive enterprise." Being a reactive enterprise is not productive, as the organization will constantly be reporting on how it recovered from a problem within its network. This tends to give the IT department a negative image. The main issue with being a reactive enterprise is that customers will be unable to get to their business application, service or assistance they need during an outage, and therefore the business is not profiting.

This is why it is essential to budget for necessary, ongoing maintenance when designing a network. This concept is no different than getting an oil change for a car every 3,000 miles or going to the dentist every six months.

Organizations should have capabilities to monitor per protocol, understand the network trends and locate network problems quickly. There are many tools, but finding one that is best for a particular organization's needs is most important. Once there is a system that can monitor the network, reporting on network health to management should take place on a regular basis. This will help justify changes, upgrades and, most important, provide data on how IT dollars are utilized.

Hierarchical Internetworking Model

Every day an enterprise relies on communication with customers and workers to make a profit. This communication is dependent on a solid network foundation; it is the backbone of an organization's success. Enterprise networks consist of many devices working together to

allow seamless communication. When making changes to the enterprise or designing new capabilities, it is important to follow architectural best practices.

Enterprise networks are complicated. There are many devices that make up the infrastructure, many different devices and hardware working together to create seamless communication. To insure highest reliability and best performance, the network foundation must be capable of many features and protocols. A solid foundation will allow the organization to grow seamlessly going forward.

When setting this foundation, the enterprise needs to understand ongoing business trends and new technologies that may change the way it does business. Having a solid foundation that can quickly react to new trends and the ability to add new technologies well positions an organization to take on more business.

Keep in mind that supporting an enterprise communication network can be expensive. When selecting devices, it is important to understand what is needed today while having the vision to understand how the organization can build from its foundation to stay ahead of competitors in the market.

This type of foundation is no different than that of a house. Throughout the years, the owner may update the bathroom or get new flooring, and the foundation is built to allow for these types of changes. With networking, services may be added or upgraded, and the foundation needs to be designed to allow for these adjustments.

As an example, many organizations have added the transport of voice services over the IP network. With a well-designed network, a change on the edge device in the access layer is all that is needed to add this service. This example highlights a seamless reaction to a new business trend.

The pieces of hardware that form the foundation of a network are basically the network switches, routers and infrastructure cabling.

Enterprise communication networks have building blocks, also referred to as layers, each of which has unique functions. Devices have unique roles and different requirements for each layer. It is possible to combine roles in a single device if the business requirements are met. These layers comprise the overall building blocks of the network.

While every network is different and has to be crafted to the needs and scale of the enterprise, there is a hierarchical internetworking model to help simplify the task of building a reliable and scalable network. This hierarchical model can help to design, deploy and maintain a scalable, trustworthy, cost-effective network. Here are the primary layers:

ACCESS LAYER: This layer includes hubs and switches. It is also called the "user access layer" or "edge" because it focuses on connecting end-user devices, such as desktop and notebook computers, printers and Voice over IP (VoIP) handsets, to enterprise networks. This layer ensures that packets are delivered to an end user's device. Reliable end-user access to business applications using either a wired or wireless connection occurs at this layer.

DISTRIBUTION LAYER: This layer includes local-area network (LAN)-based routers and layer 3 switches. This is where packets are properly routed between subnets and virtual LANs (VLANs) in an enterprise. This is also where traffic engineering is done to manipulate data flows using quality of service (QoS).

CORE LAYER: This layer is considered the wide area network (WAN) backbone of the enterprise, where high-end network devices are located. This layer does not route traffic at the LAN but moves traffic to and from the WAN. The WAN provides connectivity for remote employees and customers. In addition, no packet manipulation is done in this layer. Rather, this layer is concerned with speed and ensures reliable delivery of packets. It's the area where data moves as fast as possible with minimum delay.

(There are also small- to medium-size enterprises using the "collapsed distribution/core layer." This is where one device does both distribution and core layer functions. This is referred to as a "collapsed core architecture.")

DATA CENTER LAYER: Here is where the critical business data and applications are stored. It is an enterprise's intellectual data, its crown jewels. It is this area that, in the past four years, has seen the most advances, enabling enterprises to move data faster and utilize infrastructure more efficiently with less space and cooling requirements. The biggest driver has been the evolution of servers virtualizing applications and the consolidation of voice, video, storage and data on the same medium.

Designing for Maximum Benefit

When building communication infrastructures or foundations, there are two basic devices found in every design: the network router and switch. These devices take many forms, and sometimes both functions are in the same device. Oftentimes, these devices are virtualized. The router and switch provide the speeds and feeds for all network communications. They facilitate the core function of a network and are a necessary and important consideration in the design phase.

Network Switch: Layer 2

Switches have replaced network hubs. Switches examine each packet and process it accordingly, whereas hubs simply repeat the signal to all ports. The interworking of the switches' architecture maps the Ethernet addresses of all the devices residing on its network port and creates the content addressable memory (CAM) table. This is a dynamic table that maps the Media Access Control (MAC) addresses of the end device to ports on the switch.

A network hub does not have a CAM table, so when a packet is received on a port, it is echoed back out to all other ports. Switches only emit a packet on the port where the destination network device resides. This allows switches to communicate between connected stations at high speed, regardless of how many devices are connected to the switch. Switches also examine each packet and have the ability to echo back out to all other ports if the packet is coded correctly as a multicast or broadcast packet.

Store-and-forward and Cut-through Architecture

End-to-end application latency requirements should be the main criteria for determining LAN switch selection. LAN switches come in two basic architectures: store-and-forward and cut-through. The two switching methodologies, coupled with an overall assessment, are critical to overall success.

Both store-and-forward and cut-through are the Layer 2 switches' basic techniques for forwarding data packets. These switches also learn MAC addresses as they examine the source MAC fields of packets when stations communicate with other nodes on the network.

A store-and-forward switch will copy the entire frame into its buffer and compute the cyclic redundancy check (CRC) before forwarding the information to its destination. Since this architecture copies the entire frame, latency varies with frame length. It will take more time and processing to examine the entire packet.

The reason the cut-through switching mode is faster is because it only looks at the destination address before forwarding. Both will look up the destination address in their CAM tables and send the frame to the appropriate interface. The drawback to the cut-through switching method is that an enterprise will not catch certain packet errors that can potentially be propagated throughout the network. Many switches are available with the ability to mix both cut-through and store-and-forward architectures.

Blocking vs. Nonblocking Switches: Protocol and Architecture

In networking, terms or acronyms are often reused, such as blocking and nonblocking. Here, these terms refer to the protocols used with the switch or the throughput of the switch.

These protocols, blocking and nonblocking, refer to the links in either the distribution layer or links in the data center. It is very common to have two or more connections from one switch to another switch. The protocol used will determine if the links are blocking or nonblocking. The goal here is to prevent network loops. There are many protocols and techniques available. It is highly desirable to use a protocol that allows nonblocking.

The most common blocking protocol is Spanning Tree Protocol (STP). This protocol creates a loop-free broadcast network (commonly referred to as the spanning tree) within a network of connected Layer 2 Ethernet switches.

When there are two or more links connected to a switch, the spanning tree allows a single active path between any two network nodes. The other links are disabled and no data will pass through them. The spanning-tree protocol algorithm exchanges bridge protocol data units (BPDU) messages with other switches to detect loops, and then removes the loop by shutting down selected interfaces. This algorithm guarantees that there is one (and only one) active path between two network devices.

Virtual Switching System (VSS), Transparent Interconnection of Lots of Links (TRILL), virtual Port Channel (vPC) and FabricPath are some of the protocols that enable switches to have all links active in a nonblocking state to other switches. Rather than drill down into each of these protocols here, the focus will stay on the ability to utilize all connected links.

When discussing switches, an organization will need to decide between a blocking or nonblocking switch. This refers to the architecture of the switch and indicates whether or not a switch is capable of handling the total bandwidth of all the ports.

If the switch has a shared bus or the switching components cannot handle the theoretical total of all ports on the switch, it is considered a blocking switch. On the other hand, a nonblocking switch can handle the total bandwidth of all its ports. There is an added cost for nonblocking switches, but for almost all applications, a blocking switch with a acceptable throughput level will work just fine.

Network Router: Layer 3

Network routers are also referred to as gateway devices. They take the Layer 3 IP address and forward this information to a different IP subnet. Imagine that someone wants to send a letter from Maine to California. To do this, the letter writer must put the physical address and the state on the envelope. Otherwise, the letter will never arrive at its destination.

In networking, the IP address is used to move data from one subnet to another. IP addresses are broken into two parts: the IP number (the physical address) and the mask (the state). In Layer 2 switching, an organization could only communicate within its IP subnet and must use a router function to forward data from one IP subnet to another. This Layer 3 forwarding modifies the contents of every data packet that is sent. This process is Layer 3 routing.

Routers are responsible for interconnecting networks and selecting the best path for a packet to travel. This is done with routing protocols. These protocols determine what routes or destinations are available on the enterprise network, build their routing tables and make routing decisions.

The most common routing protocols are: Interior Gateway Routing Protocol (IGRP), Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), Intermediate System to Intermediate System (IS-IS) and Border Gateway Protocol (BGP). Not every packet needs to be routed, but it is good to have the capability in order to communicate with other networks.

Networks have many business applications, chief among them being delivering audio and video over IP networks. These frames in the network must provide secure, predictable, measurable and (sometimes) guaranteed services. Routers provide QoS by managing the delay, delay variation (jitter), bandwidth and packet loss parameters on a network. This is the key to a successful, reliable end-to-end solution. Normally, this function is found in the distribution layer of an enterprise network.

A hybrid device is the latest improvement in internetworking technology. Combining the packet handling of routers and the speed of switching, these multilayer switches operate on both Layer 2 and Layer 3. The performance of this class of switches is aimed at the core or distribution layers of large enterprise networks. Sometimes these are known as routing switches, or IP switches.

Multilayer switches look for common traffic flows, and switch those flows on the hardware-embedded application-specific integrated circuits (ASICs). For traffic outside the normal flows, the multilayer switch uses routing functions. This keeps the higher overhead of routing functions only where it is needed. Many networks utilize high-end multilayer switches.

WAN Optimization

WAN optimizers are also known as application acceleration devices. The core function of a WAN optimizer is to solve the problem of small increases in network latency that can dramatically affect the performance of applications. Another use for WAN optimizers occurs when application traffic on the WAN becomes chatty, meaning that there are a large number of small packets traversing the WAN.

There is a point where adding more bandwidth will not improve WAN performance in a significant way. In these situations, it is an issue of physics, because the information on the WAN cannot travel faster than the speed of light. A long communication distance between sites causes delay, and operational processes are not as efficient as they could be. A WAN optimizer can correct this problem. These devices manage WAN traffic by optimizing the response times of critical applications.

Convergence in the Network

Network convergence refers to storage, data, multimedia and voice applications on a single, high-performance and efficient IP network. A number of benefits occur when network convergence is used, including increased productivity and flexibility.

Network convergence allows employees and customers to connect and communicate with each other seamlessly, reliably and securely. One particular driver is the difficulty in justifying a separate network for storage traffic in data centers as a way of ensuring performance and reliability.

Some switch manufacturers have introduced universal ports that allow the enterprise to configure any one switch port to support many network transports, such as Ethernet, Fiber Channel (FC) or Fiber Channel over Ethernet (FCoE). Many organizations are turning away from having a separate switch in the data center for FC traffic and another switch for Ethernet traffic. The benefit of not supporting parallel communication infrastructures will lower overall costs in equipment, cooling and space.

Many organizations have already added their multimedia and voice applications to the data transport network. The latest trend is to move the storage network to the data network. The servers are located in the data center, with a connection to two switches: one connection to the data transport switch and the other to the storage switch. Convergence of the two networks reduces equipment, power, space and cooling requirements in a data center.

FCoE, in particular, is a protocol that has led to convergence of both storage and data on the same wire. FCoE is a storage protocol that enables Fiber Channel communications to run directly over Ethernet. This protocol makes it possible to move FC traffic across existing high-speed Ethernet infrastructure and converges storage and IP protocols onto a single cable transport and interface.

FCoE uses a loss-less Ethernet fabric and its own frame format. It retains Fiber Channel's device communications but substitute's high-speed Ethernet links for FC links between the devices. The protocol consolidates input/output (I/O) and reduces switch complexity. It also allows cutting back on the number of cables and interface cards. The adoption of FCoE has been slow, mostly because of reluctance on the part of many organizations to change the way they implement and manage their networks.

Design Consideration for High Availability

Today, networks are running seven days a week, 24 hours a day, and most enterprises have a low tolerance for outages that stop communications. This is why it is important to design networks with high availability and full redundancy. The level of availability is dependent on the organization's tolerance for lost communication, which is typically as minimal as possible.

Some of the simplest ways to improve network availability are to build it into network devices. In-service software upgrade (ISSU) support on switches and routers will provide high availability and minimize planned downtime. This feature allows an organization to upgrade software or apply bug fixes and deploy new features and services through in-service upgrade, without creating an outage for users.

Building component redundancy not only improves availability, it often allows for hot swapping or replacing components when they have failed without shutting down the system. The components that most frequently fail are power supplies and fans.

A Paradigm Shift in Network Design

Computer networks are relatively new; LANs didn't become common in most organizations until the mid-1990s. As communication technologies matured, the way networks were created also evolved. At one time, the Token Ring protocol dominated LAN transport then Ethernet came along. WAN's Multiprotocol Label Switching (MPLS) displaced frame relay networks, and virtualization technology now allows enterprises to take advantage of idle resources in data centers. Each new technological advance has an effect on network design.

Data Center Mobility

As organizations become more virtualized with their applications and network devices, it is imperative to ensure high availability to users. Users are increasingly mobile and want increasingly flexible work schedules, leaving few windows for network administrators to make changes, upgrade software, apply bug fixes and deploy new features. The traditional passive secondary data center is not going to allow enterprises to reach this goal.

At the same time, senior management questions the wisdom of leaving the resources in the secondary data center sitting idle, as well as the cost of this approach. Organizations must introduce technology that allows applications and network devices to become stateless — the ability to move an application or workload from one location to another.

With stateless computing, users typically don't notice any changes in the work that they are doing on the network. This technology can allow the organization to keep its infrastructure up to date and more effectively utilize resources in both data centers.

The challenge in networking is the link layer, the Layer 2 domain (also know as the broadcast domain). Fortunately, there is a protocol, Overlay Transport Virtualization (OTV), that allows organizations to have the same Layer 2 domain in multiple locations.

This is only half the solution, because Layer 3 traffic must also be redirected so remote users can find where the active application was moved to. There are a couple of options for achieving this.

An enterprise can utilize a device that redirects Doman Name Services (DNS) or Locator/ID Separation Protocol (LISP) traffic. LISP is a network architecture and set of protocols that implement a new semantic for IP addressing. LISP separates the "where" and the "who" in networking and uses a mapping system to couple the location and identifier.

Software-defined networking (SDN)

Existing network devices, switches and routers generally have a data plane (the part of the device that forwards packets in one interface and out another) and a control plane (the part of the device that builds the rules for forwarding packets). These functions are in the same device.

For example, a router receives a packet, consults the routing table and then forwards the packet. That's the data plane. But at the same time, the control plane is running routing algorithms, such as the BGP and OSPF protocols, and that part of the router builds the routing table in the data plane. Control planes include routing, access controls and firewalling, and traffic engineering, among other things.

With SDN architecture, the control and data planes are decoupled, network intelligence and state are logically centralized, and the underlying network infrastructure is abstracted from the applications and network services. Essentially, SDN moves the network control and decision-making away from the switching devices and into a separate network element: the flow controller. The goals of SDN are to simplify network management, enable network reconfiguration and enhance understanding of what's actually happening on the network.

The problem that SDN attempts to resolve is moving the control plane off each device and into an all-seeing and all-knowing controller that can make globally correct decisions for routing, access control and traffic engineering — better than any individual device could do. Researchers are still studying whether SDN can be used to build better, smarter and more flexible networks. For now, it's worth keeping an eye on.

The Road Ahead

This paper is a brief overview and good starting point when looking to improve an enterprise communication network. We have discussed many topics at a very high level and there were many topics that were not covered. The goal of this paper is to help guide you through the decision process when making changes to your infrastructure.

Typically, many IT professionals work in siloes and have little understanding of the overall enterprise picture, so it is important to have a vision of what the goals are before making changes. Networking professionals make changes that affect all enterprise communication. Therefore, it is critical to an organization's success to architect infrastructures with reliability and flexibility.

It is important for networking professionals to understand their environment, technology trends and their business. There are many resources available to educate, and leveraging your partners is one of the easiest ways to stay ahead. Partnering with an independent provider is a no-cost way to have a trusted advisor on your staff working to ensure your organization's success.

Video Channel

Networking

Review networking case studies and solution architect discussions: **CDW.com/routerswitch1**

CDW: A Networking Partner That Gets IT

CDW offers a wide selection of network solutions designed to increase the speed of access to critical applications and information across the enterprise. Routers and switches are the foundation for all enterprise communication. These devices are the core infrastructure that enables smooth and reliable communication. Well-architected infrastructures are equipped to take advantage of today's networking requirements and are poised for tomorrow's demands.

Your CDW account manager and solution architects are ready to assist with every phase of choosing and leveraging the right network solution for your IT environment.

Our approach includes:

- An initial discovery session to understand your goals, requirements and budget
- An assessment review of your existing environment and definition of project requirements
- Detailed manufacturer evaluations, recommendations, future environment design and proof of concept
- Procurement, configuration and deployment of the final solution
- Ongoing product lifecycle support
- Availability of on- and off-premises network managed services

To learn more about CDW's network solutions, contact your CDW account manager, call 800.800.4239 or visit CDW.com/network.

The HP Converged Infrastructure blueprint and the HP FlexNetwork architecture are designed to provide your organization with the predictable, solid performance, high availability and comprehensive management needed to support critical applications and overarching business needs.

CDW.com/hp

Juniper's EX Series Ethernet Switches address the access, aggregation and core layer switching requirements of micro branch, branch office, campus and data center environments, providing a foundation for the fast, secure and reliable delivery of applications that support strategic business processes.

CDW.com/juniper

Brocade® switches are the foundation for high-performance connectivity in storage, IP and converged network environments. These highly reliable, scalable and available switches are designed for a wide range of environments — enabling a low TCO and fast ROI.

CDW.com/brocade

A key component of the Cisco® Self-Defending Network, the Cisco Integrated Services Router allows organizations to synchronize routing and security policies and reduce operational costs while raising the level of security throughout the network.

CDW.com/cisco

SHARE THIS WHITE PAPER

CDW PEOPLE WHO GET IT