Lippis Report
White Paper

Navigating Network Infrastructure Expenditures During Business Transformations

by
Nicholas John Lippis III
President, Lippis Consulting

July 2009

Executive Summary:

This paper explores tools and techniques available to business and IT leaders who seek to maintain and increase network availability through management, device feature exploitation and network design, especially during business process change. The mixed network vendor approach to diversity and redundancy is explored and brought into question as some IT leaders, pressured by lower capital budgets, seek to procure infrastructure from low cost providers as a means to make ends meet. The paper takes the position that a common network based upon mixed network supplier platforms paradoxically reduces network availability by increasing complexity and operational cost, the highest cost component in total cost of ownership (TCO). The paper further identifies that a mixed network vendor environment restricts design options, increases security vulnerabilities and limits the ability to optimize application performance.

While the economic downturn and subsequent business transformation opportunity highlights the value of a single-vendor strategy, this approach is independent of economic cycles. For most IT leaders, during growing economies IT spending increases as does service delivery, yet during down cycles IT budgets are reduced but service delivery requirements increase due to the necessity for business transformation. Vendor selection strategy allows IT leaders to best manage this balance.

By standardizing on strategic vendors since its inception in 1999, JetBlue has been able to manage the balance between IT service delivery and spending levels over multiple economic cycles better than its competitors. JetBlue realized early on that there are material customer benefits to strategic purchasing and thus implemented this strategy from the aircraft it purchased to the IT systems it deployed. This foresight gave JetBlue a significant competitive advantage by tapping a geographically distributed employee pool of contact center agents, lower IT operational cost, and lower aircraft maintenance cost, etc. JetBlue now has a market capitalization twice that of United Airlines and nearly equal to American Airlines. JetBlue and others suggest that standardized strategic vendor relationships are a major contributor to continually and successfully managing this balance, independent of where and how fast balance points shift and allows them to respond to market changes faster than competitors.

With this in mind and based upon the discussion below the following recommendations are offered for consideration:

- Consider networking suppliers with the financial stamina to not only withstand periods of economic downturn but also enjoy increased market share and customer scenarios to guide its research and development investments to assist their customers in transforming their businesses.
- To reduce risk of network downtime and operational spend consider a single strategic network platform partner rather than a multi-vendor solution.
- Avoid procuring perceived low cost products for low functionality places in the network such as the edge as this tends to increase operational cost, the largest cost component in TCO.
- Consider selecting a network platform supplier who possesses an architectural view of high availability. Add additional weight in the vendor selection process to a supplier who conducts large-scale deployment validation and testing to achieve higher availability for an entire network.
- Consider the single network platform approach for mission critical network applications such as deployments for campus, data center, Internet edge, etc.
- Consider equipment sparing, device high availability features and network design options as components to deliver a high availability network.

Table of Contents

High Availability Networking Through Dual Backbones

As networks have become a business platform supporting most if not all business processes, business and IT leaders have become increasingly risk adverse to network downtime and rightfully so. A corporate network is the only horizontal IT asset upon which all other IT assets rely to deliver their value. The cost of network downtime is much greater than a single IT application, as a poorly designed network possesses the risk of shutting down all IT applications and thus business process. The table at right, compiled by Contingency Planning Research illustrates the high cost of downtime per industry per hour.

Industry	Application	Average Cost / HR of Downtime (US$)
Financial	Brokerage Operations	$ 7,840,000
Financial	Credit Card Sales	$ 3,160,000
Media	Pay-per-view	$ 183,000
Retail	Home Shopping (TV)	$ 137,000
Retail	Catalog Sales	$ 109,000
Transportation	Airline Reservations	$ 108,000
Entertainment	Tele-ticket Sales	$ 83,000
Shipping	Package Shipping	$ 34,000
Financial	ATM Fees	$ 18,000

With the financial cost of downtime so high the financial services industry, in particular, has a long history of deploying risk management techniques to mitigate downtime and catastrophic events. In addition to downtime cost a corporation tarnishes its "brand goodwill" along with lost productivity and revenues by lost customer contact and support during downtimes. But it's not only the financial services industries that have deployed risk management techniques to increase network availability. Many industry segments are deploying the techniques outlined below to increase network availability by reducing mean time to repair (MTTR), managing network outage risk and in the process improving business continuity and disaster planning.

One of the best examples of increasing network availability is the over design of backbone networks, commonly referred to as dual backbone. See Figure 1 on the next page. The dual backbone network is common in large corporations and provides high network availability operation thanks to redundancy at all network tiers. But as capital spending has become constrained some business leaders and risk management executives are exploring the concept of mixing network equipment vendors in a corporate network as a way to stretch capital dollars and add another degree of redundancy in the hopes of increasing availability even further.

By incorporating a dual network vendor design, the hope is that network outages due to specific vendor equipment faults and/or exploits targeting a vendor's software will be mitigated by diversifying the number of suppliers in the network. From a budgeting and risk mitigation point of view IT leaders commonly look toward the edge of their networks to deploy low cost network providers. While on the surface this approach seems alluring, deeper inspection shows that network availability actually decreases while operational cost increases in mixed vendor environments. We'll use the dual backbone network example to make our points.

Mixed Vendor Networks Drive Up Operational Cost While Reducing Network Availability

Dual backbone networks are often deployed in mission critical environments where network downtime results in significant and material consequences. In these environments risk management executives are focused on mitigating operational downtime risk while IT executives seek to stretch capital dollars. For corporate networking this translates into the following eight risk management goals:

1. High availability
2. High reliability
3. Low Mean Time To Repair (MTTR)
4. Maintaining business continuity during disaster recovery
5. Avoiding vendor lock-in
6. Achieving more favorable acquisition pricing due to competition
7. Avoiding single-vendor risks such as targeted exploits against a supplier's equipment
8. Avoiding the winding down or change of ownership of a supplier's business operations

Dual sourcing network infrastructure cannot achieve the above goals as this strategy results in fundamental disadvantages. These disadvantages include:

1. Complexity Inflation which drives up operational spend
2. Logical Networking Vulnerabilities
3. Network Services Relegated to Lowest Common Denominator
4. Lower Availability and Reliability

Complexity Inflation

Multiple network equipment suppliers within a network increase network complexity, which drives up operational cost as well as the probability of outages. From a practical point of view, operational staff who have standardized on a smaller number of vendor's management software are more proficient in its configuration, management, monitoring, troubleshooting and overall operations. When operational staff is required to support multiple vendors offering similar equipment the operational budget experiences complexity inflation or operational budgets are more stressed from the challenges of supporting multiple vendors. Complexity inflation increases human capital cost as operational staff is required to be trained and become proficient in new management software and its nuances. Complexity inflation is measured by either increased operational hours or additional staff required to manage the network. For many organizations complexity inflation is felt by operational staff being overwhelmed by an increase in tasks and assignments beyond proper workload. Unchecked complexity inflation results in network outages, delayed projects and/or re-sizing of operational staff. The end result of a dual vendor strategy is that IT may not be able to operationally support either backbone network appropriately unless an infusion of human capital is appropriated.

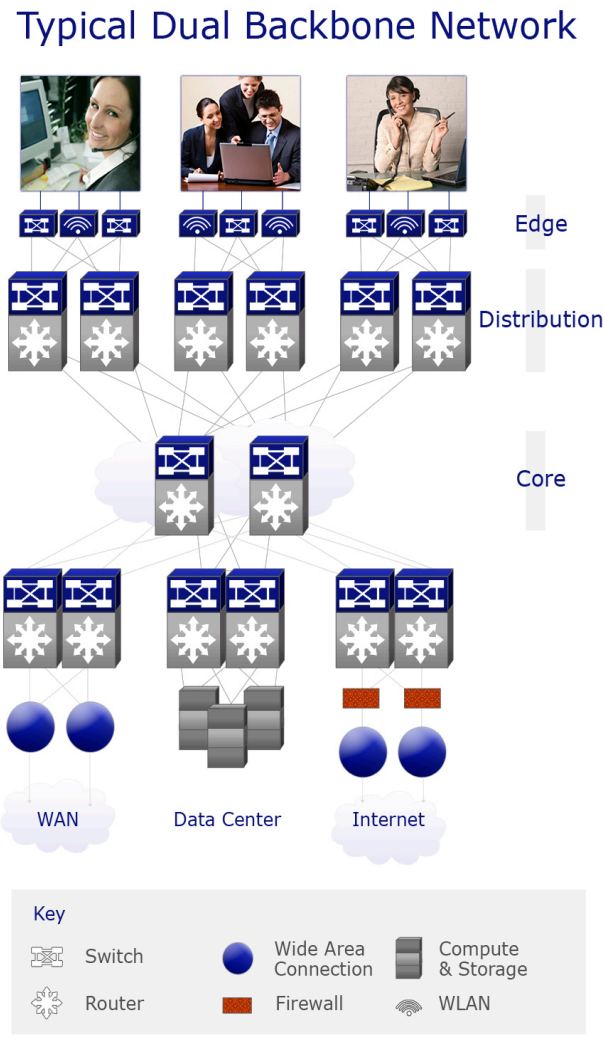

Figure 1

From a total cost of ownership (TCO) point of view, operational cost dominates acquisition and facilities cost. Typically acquisition cost represents only 25% of TCO while operations consumes 40%, leaving 35% to facilities cost. Facilities cost include wide area network service provider charges, power and space consumption plus equipment maintenance. Operational cost is dominated by human capital cost. Therefore, mixed network vendor environments drive up the most expensive cost component in TCO, that being operational cost. Further, low cost network vendor's equipment is usually deployed at the lowest cost and functionality part of the network that being the edge. This strategy increases the highest cost component of TCO, operational spend, in an effort to save few capital dollars. In short, this practice is penny wise but pound foolish. But beyond cost, mixed vendor network environments increase complexity and complexity is not reliability's best friend. In the following 3-minute Lippis Report podcast, Nick Lippis discusses network complexity and the disruptive outcomes it creates via two examples: the US Customs and Border Protection Agency at Los Angeles Airport and global Skype VoIP service.

Network Complexity Podcast

Listen to the Podcast

Logical Networking Vulnerabilities

Network equipment does not operate in isolation as it shares physical and logical connections, which influence system behavior. From a physical connection point of view, the dual vendor network strategy has value in the fact that there are redundant systems; however network devices share information and files plus execute common protocols, which they rely upon to perform their basic task of packet forwarding. For example, a malicious attack on a routing table of one vendor would result in corruption of both vendors' routing tables as routing tables are updated and shared between vendors. While business or IT executives may have hoped that the dual vendor strategy would reduce risk, what has occurred is increased complexity of problem isolation. In short, vulnerability and risk has increased.

Network Services Relegated to Lowest Common Denominator

As computer networks become more complex the "Law of the Weakest Link Always Prevails." Whenever a flaw or weakness allows a problem to occur it compromises the entire system, just as one weak section of a levee can inundate an entire community. The interconnection between mixed vendor network equipment will be via standard interfaces or in some cases the weakest link available. The networking industry is perhaps the most standard of all IT segments thanks to the IETF and IEEE organizations and the advent of TCP/IP, the Internet and local area networking.

In a dual backbone architecture deployed with dual network equipment vendors, the services offered by the dual backbone are relegated to the least common denominator of standard offerings. While there may be a compelling innovation offered by one of the vendors that an IT organization would like to implement, it can not unless those innovations are available on both platforms. The likelihood of it being available simultaneously so that IT can implement and take advantage of it is doubtful, as competitors never deliver features and innovation in unison. For example, in a mixed vendor dual backbone architecture, IT operations is relegated to deploying the least sophisticated Quality of Service (QoS) architecture despite the secondary backbone being equipped with a more sophisticated set of capabilities. Further, vendors often implement or interpret standards differently creating deployment frustration and interoperability problems. In a mixed vendor environment, the networked system must operate under the law of the weakest link.

The mixed vendor dual backbone is relegated to delivering the lowest common denominator of standardized services. Standardized services are nearly always less sophisticated and in the areas of quality of service, application intelligence and acceleration plus route optimization and network virtualization, the mixed vendor strategy limits IT optimization of availability and application performance. This set of attributes is a paradox as dual backbones are acquired to increase network availability and performance yet when mixed vendors are incorporated the direct opposite occurs. In an industry dominated by standard interfaces, service providers for example choose to build their infrastructure with a small number of network vendors to leverage innovation.

Lower Availability and Reliability

Defending against malware and exploits in a mixed vendor network poses yet another exposure and vulnerability. Network security offers defenses against crippling exploits before they propagate throughout a network infecting end-points and servers, which send operational staff into a reactionary mode to contain the exploit and cleanse compromised systems. The industry introduction of network access control (NAC) is a defensive technology to mitigate risks during and post-network access. Many NAC components are not

standardized yet and thus vendors differ on their implementations. In mixed network environments, NAC deployments would be difficult at best, resulting in islands of trusted and non-trusted networks.

In addition to NAC, network equipment suppliers have increased their response to known exploit signatures by alerting customers through security advisories and offering signature defenses to mitigate the exploit. Some IT and risk management executives find comfort in thinking that a mixed vendor network environment provides protection from exploits targeting a vendor's architecture. The thinking here is that by deploying two network vendors their risk of such an attack is mitigated. The fact is that those who design exploits seek to maximize harm while minimizing effort. To meet that end, most network attacks focus across the implementation of a particular protocol rather than a specific vendor's architecture vulnerabilities.

For example, Simple Network Management Protocol (SNMP) has come under attack thanks to its default community strings which have allowed an attacker to gain information about a device using the read community string "public", allowing the attacker to change a system's configuration using the write community string "private". The opportunity for this exploit is increased as the SNMP agent is often installed on a system by default without the administrator's knowledge.

The H.323 protocol is most commonly used in Voice over IP (VoIP) and video conferencing applications for the exchange of voice and video communications over networked systems. H.323 protocols are used across a wide range of vendors and when attacked impacts all such equipment. Some H.323 vulnerabilities are caused due to various errors in the processing of H.225 messages over TCP. This vulnerability can be exploited by malicious people to crash or reboot an affected device by sending specially crafted messages to it via default port 1720/tcp.

Both SNMP and H.323 vulnerabilities are industry-wide rather than vendor specific. Therefore, the dual vendor strategy does not protect an organization against industry-wide vulnerabilities as most networking products, at their core, are based around standard implementations of protocols which themselves can suffer from these industry-wide vulnerabilities. When working with multiple vendors on a common network infrastructure IT needs to evaluate the vendor on how they respond to these vulnerabilities with an understanding that there will be serious security issues over time that are not vendor specific. In short, dual vendor implementations are not an effective network security strategy.

The Dual Vendor Strategy: A Lot of Pain with Little Gain

As mentioned above, most business and IT executives view their networks as a platform investment, which delivers future feature dividends. When business and IT leaders make a platform decision they are not just choosing a supplier but choosing a partner that has the ability, skills, research and development, financial stamina and intent to invest in their platform. The dual network vendor strategy is a hedge across two platform investments. However, since innovation between competitors is different and their delivery not synchronized, this strategy does not allow corporations to exploit each platform's innovation as network architects are limited to delivering basic interoperable standardized services.

If network architects could deploy multiple platform dependent innovations on a common network, then their corporation would benefit by, in essence, conducting a horse race between platform providers, picking and choosing which innovations to deploy. But this is not a reality of the market. Dual backbone networks constructed with more than one vendor cannot deploy different implementations of QoS, network security, application intelligence, application acceleration, network virtualization, and paradoxically innovations in high availability, which is the original dual backbone design goal. Even worse some of these innovations could simply not be deployed, as both platform providers do not support them. This limits the design options available to network architects precluding optimization of application performance and availability. In essence the dual vendor strategy handicaps an IT organization from exploiting the investment made in the two platforms while reducing research and development dollars available to each platform provider.

The dual vendor strategy forces an IT organization to sacrifice network security, QoS, application intelligence, application acceleration, network virtualization, wide area acceleration, and high availability capabilities, among others. All of these capabilities contribute to increased reliability of applications running on the network. With these capabilities precluded, a corporation runs a higher risk profile for network outage.

The mixed vendor strategy to dual network backbone design has inherent difficulties, which are counter-productive to achieving high availability and minimized operational cost. The mixed vendor approach restricts network design options, limits the use of innovation, and increases complexity, which translates into higher operational cost and longer MTTR. While the mixed vendor strategy approach may on the surface seem logical, however, underneath the surface resides more complexity that adds rather than reduces risk. These difficulties in the mixed vendor approach are not isolated to dual backbone networks but are applicable to branch office deployments, data center designs, etc. The mixed vendor strategy can result in a slippery slope of where an IT department stops in its level of redundancy. If mixed vendor dual backbones are required then is a mixed vendor dual edge layer also required? Are mixed vendor dual wiring closets and wireless LAN controllers needed? Are dual desktop operating systems needed? Does IT provision dual application infrastructures built on different supplier's platforms delivering the same application? At what point is there enough redundancy to ensure the level of availability required to satisfy risk management and at what cost? Do business leaders appropriate duplicate IT capital and operational budgets?

High Availability Dual Backbones: Single versus Multi-Vendor

The alternative to the multi-vendor solution is a single-vendor approach to dual backbones. Clearly dual backbone architecture does not require dual vendors. The dual backbone approach is an excellent strategy to deliver high availability and business continuity. All of the difficulties identified above are not represented in the singe vendor approach. Complexity is minimized, reducing operational spend and MTTR while providing a simpler approach for IT leaders to automate post-crash streamlined business processes. There are no mixed vendor logical networking vulnerabilities. Corporations are not relegated to delivering the lowest common denominator network services, but are free to exploit all the services and innovations a network platform affords. In addition a single-vendor strategy offers a more environmentally friendly approach too by the reduction of appliances and devices, which require separate power sources.

As executive management tasks IT leaders to transform their business through IT automation, many IT leaders will find themselves seeking a balance between IT expectations and capacity to deliver business transformation. IT management is evaluating its skill sets, headcount, expertise, training requirements and infrastructure to understand how big a gap exists between expectations and capacity. Based upon the post-2008 crash period IT spending pullback, the gap is large and growing. A single-vendor strategy is a contributor to closing the expectation versus capacity gap, thanks to its complexity reduction and innovation absorption benefits.

Not only is higher availability achieved with a single-vendor architecture, but the ability for an IT organization to absorb network innovation increases too. In a mixed vendor environment, IT organizations have to wait until new services are standardized. Standards always lag behind innovation. The mixed vendor approach puts a corporation at a competitive disadvantage by limiting an enterprise's ability to absorb network innovation and automate business processes.

Vendor selection importance increases as business and IT leaders seek to standardize their networks on a single supplier. Clearly not all vendors are the same. An analysis of suppliers across the industry shows that there are a certain group of vendors that strive to meet a basic feature set. There are other suppliers who are able to check the boxes of the appropriate IETF or IEEE standards, which are usually part of the Request for Proposal (RFP) process. Then there is a class of suppliers who incorporate customer operational challenges and build solutions to them in their platform. These solutions could be high availability capabilities, which stem not only from dual hardware but software innovations that allow dual hardware to perform more reliably. For example, in the event that a switch or router suffers a hardware failure, software innovation that does not corrupt the data path so that the network continues to forward traffic without interruption increases availability. Some suppliers provide a set of high availability and software services that are available in their network platform that will load balance traffic while simultaneously recovering from a catastrophic failure of one side of the network. Other platform innovations include fast reroute, network reroute, load sharing, and redundancy technologies. Another example is the ability to manage dual network devices such as LAN switches as a single device thanks to network virtualization, which reduces operational spend while increasing availability.

With a focused platform investment on a single-vendor that supports a robust high availability network strategy, an IT organization can optimize its operational capabilities thanks to lower complexity. The value of a single-vendor network business platform can be expressed in increased return on investment and lower TCO terms when compared to the multi-vendor approach. Single-vendor networks also enjoy a wider range of design

options available to network architects, which contribute to increased availability through lower MTTR and faster innovation absorption.

There are multiple risk management techniques available to achieve high reliability such as equipment sparing and device features. Network design also contributes to high availability as discussed in the dual backbone design. Redundant links and loops to connect network devices add availability as well as particular protocols that deliver redundancy such as Virtual Router Redundancy Protocol (VRRP), etc.

Recommendations

As in the 2000-2002 dotcom bust, many start-up operations and large firms also entered Chapter 11 restructuring and Chapter 7 liquidation. The same is occurring now, with Nortel Networks liquidation being the most obvious example. There is a consolidation phase occurring in the networking industry as evidenced by as Foundry Networks merging with Brocade, HP reorganizing its proCurve group into its TSG organization, IBM tightening its relationships with Brocade and Juniper Networks, Force 10 Networks' merger with Turin Networks, etc. These reorganizations put in question product priorities, research and development levels and increase IT risk. During down economic periods most IT organizations choose to procure equipment and services from independent financially secure firms who are in charge of their own destiny. The current economic period has produced an acceleration of this buying behavior, often called a "flight to safety" or "flight to value".

While the economic downturn and subsequent business transformation opportunity highlights the value of a single-vendor strategy, this approach is independent of economic cycles. For most IT leaders, during growing economies IT spending increases as does service delivery, yet during down cycles IT budgets are reduced but service delivery requirements increase thanks to the necessity for business transformation. Vendor selection strategy allows IT leaders to best manage this balance.

By standardizing on strategic vendors since its inception in 1999, JetBlue has been able to manage the balance between IT service delivery and spending levels over multiple economic cycles better than its competitors. There are customer benefits to strategic purchasing and JetBlue realized this early on and implemented this strategy from the aircraft it purchased to the IT systems it deployed. This foresight gave JetBlue a significant competitive advantage by tapping a geographically distributed employee pool of contact center agents, reduced staff IT training requirements, lower IT operational cost, lower aircraft maintenance cost, etc. JetBlue has a market capitalization twice that of United Airlines and nearly equal to American Airlines. JetBlue and others suggest that standardized strategic vendor relationships are a major contributor to continually and successfully managing this balance, independent of where and how fast balance points shift and allows them to respond to market changes faster than competitors.

With the above in mind the following recommendations are offered for consideration:

- Consider networking suppliers with the financial stamina to not only withstand periods of economic downturn but also enjoy increased market share and customer scenarios to guide its research and development investments to assist their customers in transforming their businesses.
- To reduce risk of network downtime and operational spend consider a single strategic network platform partner rather than a multi-vendor solution.
- Avoid procuring perceived low cost products for low functionality places in the network such as the edge as this tends to increase operational cost, the largest cost component in TCO.
- Consider selecting a network platform supplier who possesses an architectural view of high availability. Add additional weight in the vendor selection process to a supplier who conducts large-scale deployment validation and testing to achieve higher availability for an entire network.
- Consider the single network platform approach for mission critical network applications such as deployments for campus, data center, Internet edge, etc.
- Consider equipment sparing, device high availability features and network design options as components to deliver a high availability network.

About Nick Lippis

Nicholas J. Lippis III is a world-renowned authority on advanced IP networks, communications and their benefits to business objectives. He is the publisher of the Lippis Report, a resource for network and IT business decision leaders to which over 40,000 business and IT executive leaders subscribe. Its Lippis Report podcasts have been downloaded over 65,000 times; i-Tunes reports that listeners also download the Wall Street Journal's Money Matters, Business Week's Climbing the Ladder, The Economist and The Harvard Business Review's IdeaCast. Mr. Lippis is currently working with clients to transform their converged networks into a business platform.

He has advised numerous Global 2000 firms on network architecture, design, implementation, vendor selection and budgeting, with clients including Barclays Bank, Microsoft, Kaiser Permanente, Sprint, Worldcom, Cigitel, Cisco Systems, Nortel Networks, Lucent Technologies, 3Com, Avaya, Eastman Kodak Company, Federal Deposit Insurance Corporation (FDIC), Hughes Aerospace, Liberty Mutual, Schering-Plough, Camp Dresser McKee and many others. He works exclusively with CIOs and their direct reports. Mr. Lippis possesses a unique perspective of market forces and trends occurring within the computer networking industry derived from his experience with both supply and demand side clients.