

A Comparison of Total Costs of Ownership of 10 Gigabit Network Deployments in the Data Center

Contents:

Introduction	2
Objectives	2
Executive Summary	3
Background	6
Methodology	7
Data Center Network 1 – Separate GbE LAN and 4G FC SAN	9
Data Center Network 2 – Separate 10GBASE-T LAN and 8G FC SAN	10
Data Center Network 3 – Converged 10GbE and 4G FCoE	10
Data Center Network 4 – Separate 10GBASE-SR LAN and 8G FC SAN	11
CAPEX Analysis and Considerations	11
Key Findings	13
Cabling CAPEX	14
Network Cards CAPEX	14
Aggregation Layer Switch CAPEX	15
Top of Rack (ToR) Switch CAPEX	15
OPEX Considerations	15
Management Considerations and Risk Factors	17
Conclusions	18

Introduction

Technological advances, such as virtualization and multicore servers, open the way to a new generation of more efficient, higher performance data centers. The new developments can bring important benefits, but they also challenge organizations to make critical infrastructure decisions that will impact the future of their IT systems and ability to stay competitive. Among the most important of these decisions is how to implement 10 Gigabit networking in the data center.

Virtualization has emerged as the strongest driver for deployment of 10 Gigabit Ethernet (10GbE) in the data center. At the same time, data center specific transport protocols now make it possible to converge previously separate networks into a single 10G connection.

Using the new network technology to converge LAN and SAN connections in the data center offers advantages over proven approaches using separate networks. To decide whether converged networks, such as Cisco's Unified Fabric Architecture (UFA) is a good investment, total cost of ownership must be assessed and the advantages balanced against some disadvantages

This white paper analyzes the network architecture alternatives, associated capital costs and operational considerations for a 10 Gigabit data center network infrastructure. Throughout the paper, disparate networks will be referred to as separate while the networks that collapse LAN and SAN traffic over a single physical connection will be referred to as converged.

Objectives

- To provide decision makers with information, particularly on total cost of ownership, that will help them make the right choices when selecting 10 Gigabit infrastructures for the data center.
- To determine the most cost effective network implementation strategy to achieve short and long-term data center performance and reliability goals.

Executive Summary

It cannot be assumed that converged networks are always the right choice. It is important to analyze and compare capital expenditure and operational considerations for converged and separate networks in the context of the specific data center.

Cost of downtime, maintaining an IT security policy, and the need to accommodate unforeseen changes may be other key factors in choosing network architecture. These considerations vary greatly between organizations depending on their data center's operational goals.

For the purposes of this white paper, four different data center networks were proposed to illustrate CAPEX and operational considerations. Three of the four examples are based on the separate cabling architecture model and one is based on the converged model. 1000 servers spanning 50 cabinets were used as a baseline for all comparisons.

CAPEX components analyzed included horizontal and intra-cabinet cabling, network adapter cards, top of rack (ToR) switches, and aggregation layer switches including SAN and LAN distribution switches. Servers, cabinets and storage arrays were not calculated as they are present in both separate and converged network environments.

In total, the separate 10GBASE-T LAN and 8Gb/s FC SAN have a CAPEX of 5.7% more than the converged ToR 10GbE with 4Gb/s FC SAN on day 1. This cost disadvantage for separate networks may be offset by other factors including greater resilience, flexibility, higher aggregate bandwidth and lower cost of future network extensions. Other factors such as server, bandwidth, and application growth over these network architecture must be thoroughly examined to show true TCO.

In all cases, cabling was the smallest cost component, ranging from 7% of total CAPEX for separate 10G networks to 5% for the converged network. The reduction in cabling with converged networks gives a small cost advantage. If separate physical layer networks are initially designed a migration to converged networks could utilize the existing cable cabling infrastructure, thus extending the initial cabling investment and leave up to 100% additional cabling (converged networks use 50% fewer cables) for future server expansion resulting in no additional cabling CAPEX.

By far the largest cost in data center network installations is the electronics. This is the area where overall life cycle cost should be most carefully analyzed and assessed. The initial installation cost of separate networks is higher, but in operation as the system grows the electronics costs for a converged network can grow to offset the initial savings.

Converged networks may also offer cost savings in network administration by combining the functions of SAN and LAN groups. To benefit from this, datacenter management, processes and procedures will have to be fundamentally changed.

Figure 1 – Separate Networks Cabling Architecture

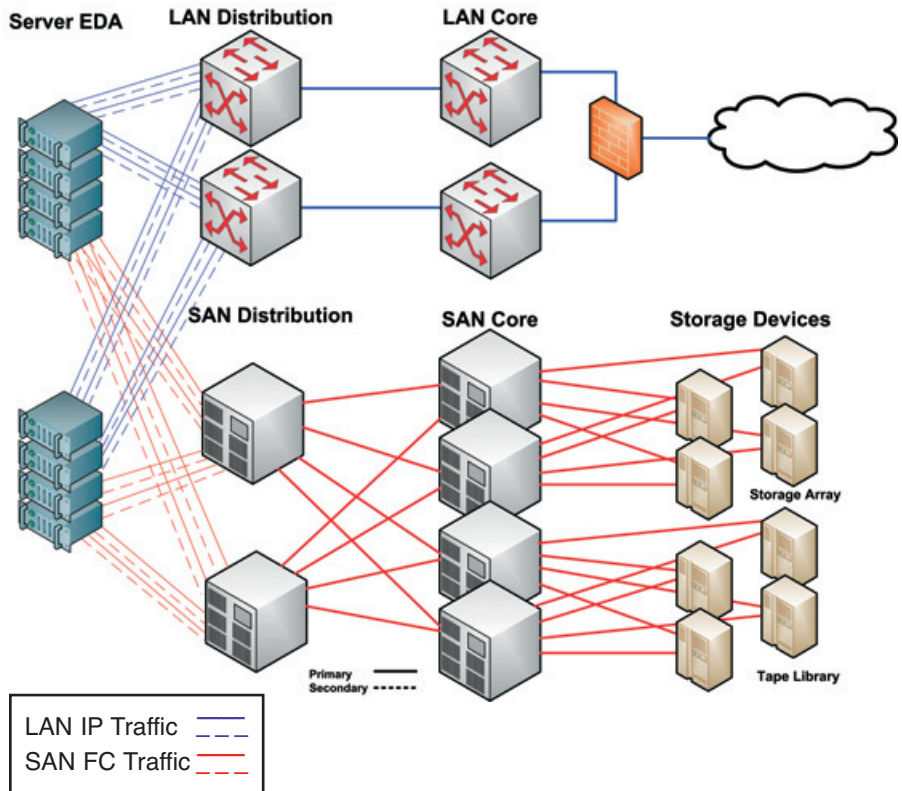


Figure 2 – Converged Network Cabling Architecture




Table 1 - Summary of Advantages and Disadvantages of Separate and Converged Networks

Cabling Architecture	Advantages	Disadvantages
Separate	Improves flexibility when re-provisioning physical connections	Cable management challenges at EDA
	Higher aggregate bandwidth	Not optimized for large/mega data centers
	Separate networks improve security	Lack of scalability can limit expansion
	Greater resilience to network failures	Re-deployment of existing cable can be time consuming effort
	Optimal switch port utilization	
	Leverages structured cabling design principles	
Converged	Simple cabling architecture	Lacks flexibility when re-provisioning physical layer
	High density server environments	Does not account for Out of Band management
	Optimized for virtualization	Combines LAN/SAN failure domains to physical switch
	I/O consolidation	First generation technology interoperability risks

Background

Electronics for 10GbE deployments have been slow to market even though suitable connectivity infrastructure has been available for several years. Category 6A cabling and laser optimized multi-mode fiber able to support 10GbE are already installed in many data centers. However, 10GbE links to the server were not, until recently, a requirement for many applications.

Now, introduction of virtualized systems is increasing the need for high bandwidth, driving demand to deploy 10GbE networks. At the same time, data center specific transport protocols are being introduced to unify disparate data center networks into a single server connection.

Separate architectures have separate networks to accommodate the different service level parameters of LAN and SAN network protocols. LANs transport Ethernet protocol and SANs transport Fiber Channel (FC) protocol.

The major difference between the two protocols is that FC is not tolerant of delays, dropped packets, or interruptions in service. This makes FC networks more tedious to design and administer. It is not unusual to see several levels of redundancy coupled with high availability at the physical layer.

Ethernet, in contrast, is designed to handle delays in service, so there are fewer restrictions on network design than with FC networks. An obvious difference in the physical layer is that the majority of Ethernet networks use twisted pair copper cabling while FC SANs use fiber optic cabling.

IP LANs and FC SANs in the data center were conceived as separate entities. Segregation ensures that service level requirements of the respective IP and FC protocols used to deliver the data between endpoints are upheld.

Introduction of new protocols, such as Converged Enhanced Ethernet (CEE) / Data Center Ethernet (DCE) and Fiber Channel over Ethernet (FCoE), has produced a paradigm shift in data center network architectures. There are fundamental networking changes affecting the OSI 7-layer stack, but the most visible difference is seen in the physical and transport layer, with the use of a ToR switch placed in the server cabinet.

This is the case when Cisco's Unified Fabric Architecture (UFA) is used to converge LAN and SAN networks. This approach uses a ToR switch within the cabinet with a single I/O cable handling traffic to and from all servers in that cabinet. The logical convergence is decoded at each server by a new type of network interface card (NIC).

Benefits of a converged network include bandwidth capable of supporting 10GbE and 4G FC, optimized for virtualized environments. Convergence ensures traffic is prioritized for movement of virtual machines between physical servers, however implementations using top-of-rack architecture may limit the servers' collective bandwidth run to the cabinet. The reduction of total bandwidth allows for the number of inter and intra-cabinet connections to be consolidated.

In addition to giving benefits, a converged network may also disrupt well established policies, practices and divisions of responsibility within the data center. Currently, most data centers are run by three separate groups, handling the network, the storage, and the applications. In a converged network environment, all three functions may use the same equipment, so separate groups will share responsibility for the same hardware.

Methodology

For this whitepaper, four different models for data center networks were created to illustrate CAPEX costs, OPEX considerations and available bandwidth. Three of the four models are based on a separate physical network cabling architecture model and one is based on converged architecture.

In each network, 1000 servers were used as a baseline for comparison of measurable CAPEX costs and operational considerations.

To eliminate any bias, the models are not based on specific technical approaches such as CEE, DCE, FCoE, or UFA. Analysis focuses on CAPEX and OPEX considerations and total cost of ownership rather than technical differences of the associated transport protocols.

During analysis of the alternative network models, differences in cabling infrastructure costs were examined alongside the cost of electronics for a next-generation data center. The application considered is server virtualization.

During analysis of the alternative network models, differences in cabling infrastructure costs were examined alongside the cost of electronics for a next-generation data center. The application considered is server virtualization where multiple operating systems and applications can run independently on a single physical server. Average ratios are 4:1, but 10:1, or 40:1 can be seen depending on the application's resource requirements.¹

¹ Building, Planning, and Operating the Next-Generation Data Center; IDC Insight. Michelle Bailey; 2008

The electronics considered include:

- Network adapter cards
- Top of rack switches
- Aggregation switches

Physical connectivity infrastructure includes:

- Multimode (MM) fiber and Unshielded Twisted Pair (UTP) copper patch panels
- MM fiber and UTP copper cabling
- MM fiber, UTP copper, and SFP+ Twinax patch cords

CAPEX costs are based on the selected vendor’s average selling price available at the time of publishing.

The bandwidths of the four networks differ. When considering any network installation it is important to understand current and future application bandwidth requirements. It should also be noted that if an upgrade to higher bandwidth electronics is required or a FCoE is implemented at the network core, existing cabling can be used; there is no need for a fork lift upgrade.

Table 2 – Examples of Data Center Network Architectures Applications and Bandwidth

	Cabling Architecture	Transport Protocols	Network Bandwidth
Data Center Network 1	Separate (direct connect)	Ethernet, FC	10/100/1000BASE-T 4G FC
Data Center Network 2	Separate (direct connect)	Ethernet, FC	10GBASE-T 8G FC
Data Center Network 3	Converged (Top of Rack)	DCE, FCoE	10GbE 4G FCoE
Data Center Network 4	Separate (direct connect)	Ethernet, FC	10GBASE-SR 8G FC

Of the 1000 servers used in each network, 20 are allocated to an EDA (Equipment Distribution Area) cabinet. All networking equipment is assigned to a single HDA (Horizontal Distribution Area).

Types and quantities of supporting infrastructure located in the EDA and HDA in each model are detailed. Networking equipment and physical layer cabling and apparatus are also detailed.

The specific system configuration for each data center network is shown below in Table 3.

Out of band (OOB) device management was not accounted for as this would be needed in all the networks. It is, however, acknowledged that there will be a difference in the cost of implementing OOB management using separate or converged networks. OOB management for ToR switches was not accounted in this model. It does not significantly change the model conclusions; however it must be accounted for during usage.

Table 3 – Separate and Converged Data Center Network Architecture Configurations

	EDA Networking Equipment	EDA Physical Layer Connectivity	HDA Networking Equipment	HDA Physical Layer Connectivity
Data Center Network 1 - Separate GbE 4G FC	(4) GbE NIC (2) 4G FC HBA	(4) 24-port Cat6A rated copper patch panels (80) Cat6A 3m patch cords (1) 48 duplex LC port pre-terminated fiber patch panel (40) Laser optimized MM fiber jumper cords	(12) Cisco Catalyst 6509E LAN switch, 48 port GbE line cards (6) Cisco MDS 9509 SAN switch, 48 port 4G FC line cards	(4000) Cat6A rated cables (167) 24-port Cat6A rated copper patch panels (4000) Cat6A 3m patch cords (100) 48 fiber MPO trunk cables (50) 48 duplex LC port pre-terminated fiber patch panels (2000) Laser optimized MM fiber jumper cords
Data Center Network 2 – Separate 10GBASE-T 8G FC	(2) 10GBASE-T Intel Pro AT Server Adapter (2) 8G FC HBA	(2) 24-port Cat6A rated copper patch panels (40) Cat6A 3m patch cords (1) 48 duplex LC port pre-terminated fiber patch panel (40) Laser optimized MM fiber jumper cords	(84) Extreme Networks Summit X650-24t switch (6) Cisco MDS 9509 SAN switch, 48 port 8G FC line cards	(2000) Cat6A rated cables (84) 24-port Cat6A rated copper patch panels (2000) Cat6A 3m patch cords (100) 48 fiber MPO trunk cables (50) 48 duplex LC port pre-terminated fiber patch panels (2000) Laser optimized MM fiber jumper cords
Data Center Network 3 - Converged 10GbE DCE 4G FCoE	(2) 10GbE CNA cards (1) Cisco Nexus 5020 ToR switch (2) 4x4 10GbE/4G FCoE expansion modules	(40) CX1 5m patch cord (1) 18 duplex LC port modular pre- terminated fiber patch panel (16) Laser optimized MM fiber jumper cords	(7) Cisco Nexus 7010 switch, 32 port 10GbE optical cards (10GbE ports only) (2) Cisco MDS 9509 SAN switch, 48 port 4G FC line card	(200) 12 fiber MPO trunk cables (17) 48 duplex LC port pre-terminated fiber patch panels (800) Laser optimized MM fiber jumper cords
Data Center Network 4 - Separate 10GBASE-SR 8G FC	(2) Intel PRO 10GbE SR Server Adapter (2) 8G FC HBA	(2) 48 duplex LC port pre-terminated fiber patch panel (80) Laser optimized MM fiber jumper cords	(18) Cisco Catalyst 6509E LAN switch, 16 port 10GbE optical line cards, X2 transceivers (6) Cisco MDS 9509 SAN switch, 48 port 8G FC line cards	(200) 48 fiber MPO trunk cables (100) 48 duplex LC port pre-terminated fiber patch panels (4000) Laser optimized MM fiber jumper cords

Data Center Network 1 – Separate GbE LAN and 4G FC SAN

This configuration represents the majority of current data center network design. It has GbE (Gigabit Ethernet) electronics, but its physical layer connectivity could support 10GbE if the electronics were upgraded to this specification.

Network 1 couples four discrete GbE copper connections for connectivity and provides two redundant LAN connections. It also has a single backup connection and a fourth connection reserved for additional LAN, backup, archiving or an iSCSI application.

Cost for the GbE NICs was estimated at \$50 per port. A 4G FC SAN was chosen because of its popularity in data centers today, which is confirmed by Dell'Oro's 2008 SAN forecast report.

The 4G FC SAN accounts for a 75% share of the total number of FC HBA ports. The cost of 4G FC HBA ports was estimated at \$325 each¹.

The horizontal connectivity chosen for this network is consistent with the recommendations of most, if not all, structured cabling manufacturers. It is also supported by independent industry research on cabling media performance specifications and referenced in the Cisco UFA white paper².

Cable runs were averaged at 100 feet for both copper and fiber. Category 6A cable is terminated to Category 6A outlets mounted in 24 port patch panels in the EDA and HDA. Laser optimized multi-mode fiber trunk cables were chosen in conjunction with pre-terminated fiber modules located in rack-mounted apparatus at both the EDA and HDA.

This network configuration provides a direct-connection cabling architecture using an inter-connect patching scheme. Standard three meter patch cords are used, providing slack for intra-cabinet cable management.

The Catalyst 6509E Ethernet LAN switch is used. This has been popular among data centers as an aggregation switch and is common in existing installations.

48 port GbE line cards were chosen to provide connectivity to each server NIC port.

Cisco MDS 9509 SAN directors were chosen with standard 48 port 1/2/4G FC line cards and populated with 4G FC SFP modules for connectivity between each servers' 4G FC HBA.

1 Dell'Oro Group SAN Five Year Forecast 2008-2012; Technology Segments SAN Switches Host Bus Adaptes (HBAs); Vol 5, No 2 S2A; July 2008

2 Cisco White Paper; Data Center Top-of-Rack Architecture Design, February 2009; "Cat6a shielded and unshielded products are designed to extend usable bandwidth up to 500 MHz and to drastically reduce alien crosstalk interference"... "10GBASE-T is currently power inefficient when compared to CX-1 and optical technologies. However, as more power-efficient third- and fourth-generation silicon architectures coupled with IEEE 802.3az Energy-Efficient Ethernet become available over the next few years, 10GBASE-T LAN on motherboard (LOM) and dense network switching products will become technically and economically viable server connectivity options."

Data Center Network 2 – Separate 10GBASE-T LAN and 8G FC SAN

This configuration consists of a 10GBASE-T LAN and an 8G FC SAN using a direct-connect cabling architecture. It provides maximum connectivity bandwidth and maintains the port utilization and flexibility to implement zone distribution or direct connect cabling architecture.

Intel's PRO 10GBE 10GBASE-T AT server adapter was chosen as the NIC card and priced at \$500 as per the average selling price. 8G FC HBA prices, estimated to be \$378 per port³, were based on Dell'Oro forecasts.

The cabling connectivity network configuration chosen is similar to Data Center Network 1. The only difference of is that two 10GBASE-T NICs per server replace four GbE ports. Total Ethernet bandwidth is increased by 5x allowing OOB management, backup or archiving, and iSCSI network interfaces to be virtually configured to run over either physical connection while providing redundant connections.

An Extreme Networks Summit X650-24t stackable 10GBASE-T switch was selected as an aggregation switch as there was no equivalent Cisco offering available at the time. It was chosen because it is a commercially available 10GBASE-T switch, optimized with front-to-back airflow consistent with accepted data center cooling strategies. In addition, the configuration chosen allows it to be stacked while managing up to 192 ports with the use of an expansion module with uplink ports.

Data Center Network 3 – Converged 10GbE and 4G FCoE

In this configuration, cost comparisons are based on Cisco's Nexus Unified Fabric Architecture. Each Nexus 5020 ToR access switch uplinks to a Nexus 7010 data center director and MDS 9509 SAN director. The ToR design reduces the amount of cabling required between HDA and EDA, but by doing so, limits the connectivity options for physical servers not located within reach of the same ToR switch. In this model, servers are connected to ToR switches spanning two adjacent cabinets as to prevent unused ports. This ensures high port utilization, but adds complexity to cable management.

Each 10GbE port requires SFP+ modules compatible with a Converged Network Adapter (CNA) server card. Two single port CNA cards provide redundancy at the server level and were priced at \$900 per card.

Each EDA cabinet houses a single Cisco Nexus 5020 ToR switch providing 10GbE DCE capability (including useable bandwidth for 4G FCoE). Five-meter (5m) CX1 patch cords were selected to provide enough length to reach a Nexus 5020 in an adjacent cabinet.

Connectivity for a converged network between EDA and HDA for Ethernet and FC traffic is fiber cabling. However, out of band management for both Nexus 5020 and application servers may require twisted pair copper horizontal cabling or the addition of a second ToR switch.

For redundancy purposes, the configuration includes the use of two 4-port 10GbE / 4-port 4G FCoE expansion cards resulting in 16 duplex LC fiber connections per switch to the HDA.

Structured cabling products chosen for this implementation take into account the emerging 40 and 100GbE standards using parallel optics as well as HA configurations. Four 12 fiber MPO trunk cables were used in conjunction with three 12 fiber pre-terminated duplex LC modules. The additional unused 12 fibers provide an upgrade path to support four 40GbE connections or two 100GbE connections as per the proposed standards.

³ Dell'Oro Group SAN Five Year Forecast 2008-2012; Technology Segments SAN Switches Host Bus Adapters (HBAs); Vol 5, No 2 S2A; July 2008

Standard one meter fiber patch cords are used to provide connectivity between ToR switches and pre-terminated fiber modules located in each server rack.

Aggregation layer switches consist of a Nexus 7010 data center director and MDS 9509 SAN director. Nexus 10GbE line cards have 32 SFP+ optical ports at a 4:1 oversubscription ratio. Uplink ports between Nexus 5000 and Nexus 7010 were chosen to be true, non-oversubscribed, 10GbE ports using every fourth (or yellow) port.

Line cards populating the MDS9509 were 4G FC since, at the time of publication, FCoE was technically stable running 4G, and 8G FC uplink modules were not available in the Nexus 5000 offering.

Data Center Network 4 – Separate 10GBASE-SR LAN and 8G FC SAN

This configuration uses a separate Cisco networking solution with an all fiber network cabling architecture. It is included for completeness as it is a viable configuration currently available to support 10GbE.

The optical 10GbE NIC server cards chosen for this data center network were Intel's PRO 10GbE SR Server Adapter (\$3200) and Cisco's 10GBASE-SR X2 Module (\$1350). Prices used for these products are ASPs (average selling prices) published by the selected vendors available at the time of publishing. The 8G FC price assumption and configuration were as described in Data Center Network 2.

Laser optimized fiber trunk cables were used in conjunction with 96-fiber, high density pre-terminated shelves allocated to both the EDA and HDA. This configuration provides a direct-connection cabling architecture using an inter-connect patching scheme.

Standard 3m patch cords are used, providing enough slack for intra-cabinet cable management.

Catalyst 6509E Ethernet LAN switches are populated with 16-port 10GbE line cards and 10GBASE-SR X2 modules. MDS 9509 SAN directors populated with 8G FC line cards and 8G SFP modules were chosen for connectivity between each server's 8G FC HBA.

CAPEX Analysis and Considerations

The following shows the four main areas where CAPEX dollars are spent in data center network deployments. Costs for servers, cabinets, storage arrays, SAN edge switches and core networking switches were not calculated as they are present in both separate and converged network environments.

Cabling CAPEX – include a physical cabling interconnect system where a patch panel port in the server EDA cabinet corresponds to a single switch port in the HDA. ASPs for standard Category 6A UTP copper and OM3 fiber horizontal cabling solutions are used for the calculation.

Compliant patch panels were chosen and distance between server and network cabinets were averaged at 100ft.

Converged networks include the cost of pre-terminated fiber cabling between HDA and EDA to be used for ToR uplinks and CX1 patch cords for downlink server connections.

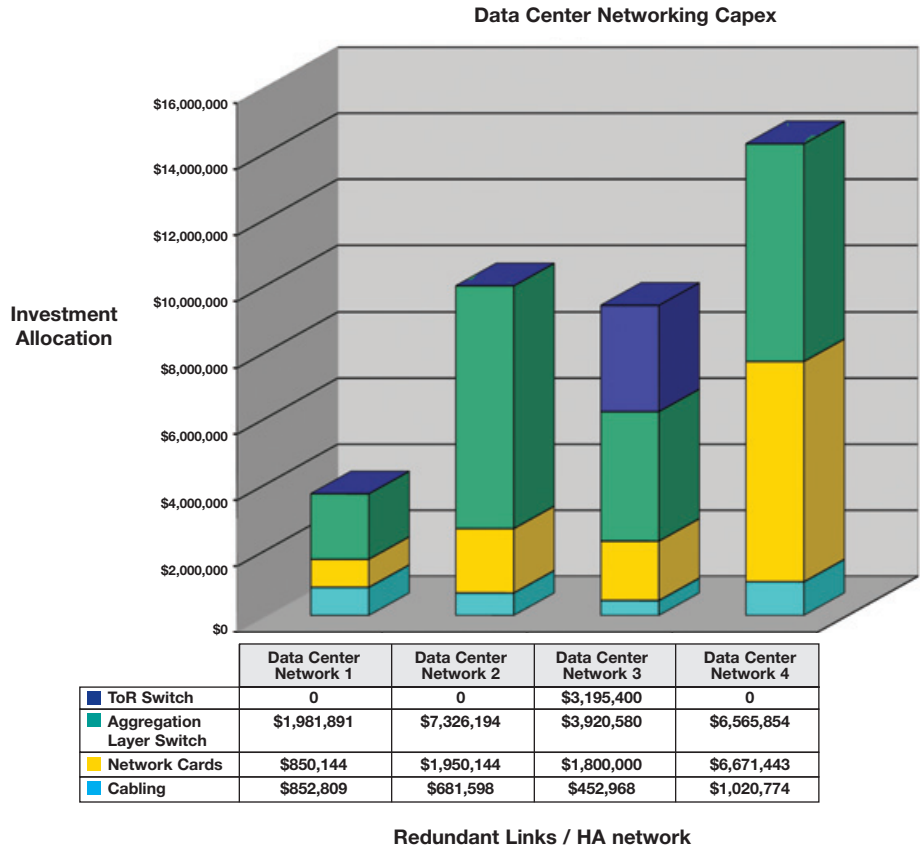
Network Adapter Card CAPEX – include the cost of NICs (network interface cards), HBAs (Host Bus Adapters), and CNAs (Converged Network Adapters). Since the market for GbE NICs and 4G FC cards is mature, ASPs from Dell'Oro's 2008 SAN Forecast are used. Prices for CNAs are based on average selling prices for a single port card. Dual cards were used to provide redundant connections. Single port 10GBASE-T NIC cards chosen were Intel's Pro 10GBE AT server adapter. Dual cards were also used to provide redundant network connections.

Aggregation Layer Switch CAPEX – these must be carefully defined for each network configuration. In the separate architecture (Data center networks 1, 2, and 4), aggregation server ports and aggregation switch ports have a 1:1 ratio. Separate networks aggregation switches account for each server's network port thus a large HDA can be expected as compared to a converged network.

A converged network uses a top of rack switch in conjunction with an aggregation switch. A port consolidation occurs using a top of rack switch. Uplink ports to each aggregation switch reduce the total number of aggregation ports and thus switches needed to support 1000 servers. Costs for licensing, standard 24x7 1 year maintenance contracts, and any additional components (power supplies and controller modules) for the switch to operate without a single point of failure were included.

Top of Rack (ToR) Switch CAPEX – include the cost of each top of rack switch, fabric and protocol licensing, and any additional components for the switch to operate with redundant points of failure. The cost of optical uplink modules are calculated using Cisco's average selling price.

Figure 3 - CAPEX Investment Allocation for the Four Data Center Networks



Data Center Networking CAPEX Investment Allocation

Key Findings

The CAPEX investment allocation in Figure 3 revealed a dramatic difference between Data Center Networks 1 and 4. The investment for Networks 2 and 3 are more closely comparable in bandwidth and CAPEX and than Networks 1 and 4.

CAPEX for a separate 10GBASE-T LAN using 8G FC SAN (Network 2) is around 5.7% higher than for a converged ToR 10GbE using 4G FCoE SAN (Network 3).

The most notable difference between the two is the allocation of CAPEX in electronics. The converged ToR 10GbE network using 4G FCoE SAN allocates approximately \$3.9M (42%) to aggregation switches and \$3.1M (34%) to ToR switches.

CAPEX investment should be segregated by electronics and infrastructure. This makes it easier to see where and if costs removed in one area are being transferred to the other.

The following were revealed from the CAPEX Investment Allocation comparison.

Cabling CAPEX

Cabling costs for separate and converged networks, each with 1000 servers show converged networks can save up to 50% of the cost of separate networks, however this initial savings in cabling CAPEX should be considered a small portion of the TCO. This is due to a reduction in horizontal cabling between HDA and EDA locations as well as fewer server network adapter ports.

The savings in horizontal cabling are offset by higher cost, direct attached, CX1 copper patch cords that connect the ToR switch to each server. Nearly two thirds (65%) of cabling costs in a converged ToR network can be attributed to the CX1 cords.

CX1 cords are a proprietary design that may be limited to current generation servers; they might not be reusable when servers are upgraded. In addition, CX1 cords may be vendor specific, and may not be compatible with servers from other vendors. This could be an important issue as extra costs for patch cable replacement may be greater than upfront savings on cabling with converged networks.

Cabling in separate networks allows for greater total aggregate bandwidth per server cabinet. This higher bandwidth allows for more servers to be added per cabinet in preparation for converged networks that are implemented at the network core in the future. Applications running 8G FC over laser optimized multimode fiber have twice the bandwidth of the 4G FCoE in a converged network today. Proposed standards that support 16G FC across the same link can be achieved and are dependent on distance and fiber specifications. Migration from separate to converged networks while adding server capacity can be achieved with no additional CAPEX for cabling infrastructure if planned for in advance.

In summary, converged networks provide a cabling CAPEX savings while separate networks have flexibility that can give OPEX benefit when adding physical servers. These factors have to be calculated and balanced for each individual installation to find where the overall cabling cost advantage lies.

Network Cards CAPEX

Expenditures on networking cards differ significantly between the four networks considered, depending on the type of cabling and protocol used.

Network 4, using separate 10GBASE-SR coupled with 8G FC network cards, requires an investment of \$6.6M. This nearly equals the \$6.5M needed for aggregation layer switches. The high cost is due, in part, to the short reach 10GbE optical transceivers and supporting network adapter cards.

Network cards for GbE coupled with 4G FC, as used in Network 1, are the least expensive. However, they provide reduced bandwidth as compared to the other three network configurations.

The converged Network 3 runs 10GbE and 4G FCoE using two single ported converged network adapters to deliver both LAN and SAN networks. This gives a 7.5% cost saving over Network 2, which supports the LAN and SAN with separate network adapters. This saving may not outweigh the flexibility and performance gains provided by separate 10GBASE-T LAN and 8G FC SAN connections of Network 2.

If server expansion is needed within a server cabinet additional ToR switches must be purchased to support the expansion, and sub-optimal port usage will exist. Separate networks allow for less CAPEX to implement server expansion while maintaining excellent port utilization of network switches.

Based on the above analysis, there are clear advantages in combining maximum flexibility, with higher aggregate bandwidth at a lower cost. Separate 10GBASE-T infrastructure with laser optimized fiber supporting 8G and 16G FC offer these benefits today. Although most current applications may not require this level of performance, virtualization applications were designed to utilize high bandwidth connections and applications teams will benefit from the increased throughput capacity.

Aggregation Layer Switch CAPEX

In addition to reducing the number of cabling channels, converged networks cut costs by reducing the port counts of the aggregation layer switches. This is achieved by using ToR switches aggregated in a HDA located where networking equipment resides.

CAPEX for optical transceivers that support the full 10GbE specification is high compared with 1G optical transceivers. However, it is offset by savings in supporting the EDA with fewer connection channels. Aggregation switch costs for separate networks vary significantly according to the chosen protocol, bandwidth and media interfaces.

Top of Rack (ToR) Switch CAPEX

Cisco's UFA solution uses ToR switches to converge LAN and SAN networks in the EDA. Separate cabling architectures used in data center networks do not require ToR switches. Instead, these architectures use passive fiber and copper patch panels requiring considerably less CAPEX than ToR switches.

For a realistic comparison, the CAPEX for aggregation layer and ToR switches in converged networks must be compared with CAPEX for aggregation layer switches and cabling in separate networks.

When comparing separate 10GBASE-T / 8G FC to converged 10GbE / 4G FCoE, the separate model costs approximately 5.7% more. However, the separate model offers 20G total LAN bandwidth and 16G SAN bandwidth. The converged model offers 20G of aggregate bandwidth, with a maximum of 8G FC bandwidth per server.

Separate models offer greater bandwidth, availability and increased flexibility when physical layer connections require moves, adds, or changes.

OPEX Considerations

In addition to CAPEX, ongoing operational cost for upkeep of a data center network must be assessed to understand total cost of ownership (TCO).

When comparing separate and converged networks, several operational factors must be considered. These include the following:

Power Usage and Cooling – Although a measurable parameter, power usage is assessed at an aggregate level. The proportion of passive versus active infrastructure components is a key factor in total data center power consumption. Adding ToR network switches in server cabinets increases their power and cooling requirements. Typically, data center cooling issues increase as server cabinet power densities grow.

A reduction in cable bulk in a raised floor pathway through use of fiber trunk cables or reduced diameter Category 6A improves air flow. This helps to prevent under floor cooling inefficiencies and allows for separate networks to be designed to support high density server deployments.

Advances in Ethernet power efficiency can offset the energy consumption of 10GBASE-T electronics. Energy saving techniques associated with 10GBASE-T networks allow up to a 65% reduction from full load power when using Wake on LAN (WoL)⁴ and Energy Efficient Ethernet (EEE) (IEEE 802.3az).

These technologies are expected to become standard offerings for LAN on Motherboards (LoMs) in the coming years, giving efficiencies without the expense of external network cards. These techniques, combined with the need for fewer high bandwidth connections, are expected to help reduce aggregate power levels, giving similar power saving benefits to a converged network.

Changes to Electronics and Underlying Infrastructure - Design, implementation and commissioning a data center can take several years. Once in production, changes to the original design may soon be needed to accommodate unforeseen events, mergers and business decisions. This requires a network architecture that combines active network components and underlying physical infrastructure in a way that makes it adaptable to changing requirements.

Moves, adds and changes (MACs) to physical layer connectivity will exist in both separate and converged networks. Converged networks using cabling that directly connects a ToR switch can make it tedious to re-provision physical networks. Re-provisioned servers are highly likely in virtualized environments and benefit from the flexibility provided using structured cabling.

Projected Growth – Growth can come in at least two forms; the need for more physical electronics and the need for higher bandwidth devices to support high-performance applications.

Although some growth can be planned and provided for, data centers often see demand growing faster than expected or face changes in business requirements. This can be a problem for data centers that employ a non-structured cabling architecture; incurring high costs to change or upgrade active equipments.

In contrast, separate networks leveraging structured cabling solutions have the ability to scale to meet growing bandwidth demands more effectively than a converged network.

⁴ Making 10GE Green – 10GBASE-T and Wake of LAN; Bill Woodruff, Data Center Journal - January 2009

To provide maximum adaptability, a network must have well chosen physical layer architecture and active components. Both converged and separate networks can be expanded by changes in either area. However, it is more crucial that converged networks using ToR switches properly provide for bandwidth growth during the original installation.

Cabling solutions used to build the network must consider evolving standards and the future directions of the data center's IT infrastructure. If they do not, there may be a risk of incompatibilities when network equipment vendors implement next generation capabilities.

Management Considerations and Risk Factors

Today's IT departments may include many separate groups; this white paper considers three that are most affected by the data center CAPEX choices. These are storage administrators, network administrators and applications groups.

With converged networks, integrating the aggregation layer switch ports for LAN and SAN connections requires a change in operational philosophy. Coordination between storage, networking and application groups becomes critical in the upkeep and maintenance of aggregation switches in a converged network.

The challenge of meshing these groups must be accepted by operators, administrators and engineers, even if running a converged network needs fewer skilled employees.

If service levels are to be maintained, IT management structures for converged networks must evolve and become less rigid. Converged network cross-functional teams will have to address a wide range of responsibilities and have an adaptable structure.

In addition, there are different risk and reward considerations related to servicing, upgrading, and maintaining aggregation layer switches in a converged network. As more end-points are aggregated, aggregation switch outages in a converged network will affect more hosts than similar outages in separate networks.

For example, a SAN switch in a separate network providing 384 ports, if taken offline, would affect performance of up to 384 servers in a high availability (HA) environment. The same 384 SAN ports in a converged network would be providing bandwidth up to 96 ToR switches affecting thousands of servers in a converged network.

Converged networks inherently increase the risk of network unavailability in the event of hardware, software or human failures. Parallel architectures offer a higher fault tolerance than converged architectures. As a result, the 10GbE/8G FC, four-connection model of Network 2 has higher fault tolerance across both networks as well as a higher aggregate bandwidth than Network 3.

For organizations where data center downtime costs thousands or millions of dollars a minute, these considerations can have huge implications. CAPEX and OPEX savings accumulated over several years could be negated by a single, brief network failure.

Data centers that support businesses with stringent uptime requirements may have to over-engineer the network to ensure that equipment failures do not violate service level obligations.

Conclusions

TCO analysis must account for both the organization's current and future needs. CAPEX investment analysis should also separate electronics and infrastructure costs. Separating these items allows the end-user to make the best decision based on their IT and cabling infrastructure refresh habits. This makes it easier to see where there are genuine savings and where investment is simply being moved from one area to another.

- **Network OPEX considerations are difficult to quantify precisely. However, the impact of different architectural choices on OPEX must be carefully considered and understood. This is essential in choosing a solution that can support all required applications with minimized total cost.**
- **Internal IT Management concerns and operational risk factors should always be taken into account. These will vary greatly from one organization to the next and although it may be hard to calculate the cost, they must be carefully considered and quantified, even if the figures arrived at are very approximate. Security policies may also not allow traffic from LAN and SAN networks to be converged over common media.**
- **The size of an organization's data center, IT operational structure and plans for growth and change and cost of downtime are all determinants in whether a separate or converged network is the right choice.**

From the perspective of performance, availability and flexibility, a separate 10GBASE-T network coupled with 8G FC SAN is best suited for today's and tomorrow's data center networks. The same physical layer infrastructure detailed in our TCO model that supports these network applications can be installed today and will support the lower cost electronics running GbE and 4G FC SAN. Migration to 10GBASE-T LAN and 8 or 16G FC SAN will then only require changes to network equipment without having to replace underlying cabling infrastructure. Users who do not require 10GbE today would be able to choose the cabling architecture outlined in Data Center Network 1 with full capability to leverage their separate networks for 10GBASE-T and 8G FC as the technology matures and prices decrease.

Even if a converged network is desired in the future, installing a separate cabling architecture now will provide greater flexibility and lower cost for adapting to present and medium-term needs.

Server upgrades and additions are well suited to be handled in a separate cabling architecture even as users migrate to FCoE. The additional cabling of a separate network can be re-used by new servers that run over converged networks without any additional cabling CAPEX. As with all new technologies, the cost of converged networks will fall as technology advances, thus delaying purchase may be the right financial decision.

Investments made today in a high quality structured cabling solution will ensure data center infrastructure can accommodate several networking and server equipment upgrades and the operational performance and reliability of the data center is upheld.



© 2009 CommScope, Inc. All rights reserved.

Visit our website at www.commscope.com or contact your local CommScope representative or BusinessPartner for more information. All trademarks identified by ® or ™ are registered trademarks or trademarks, respectively, of CommScope.

11/09 TD-A-5