ORACLE

FUSION MIDDLEWARE

DATA INTEGRATOR

An Oracle White Paper
March 2010

Ethernet: The Drive for Convergence

ORACLE

Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Table of Contents

Executive Overview

Convergence of networking technologies is a long held requirement for IT managers. The benefits gained by having a single infrastructure for data networking, storage networking, and inter-processor communications provide improved Total Cost of Ownership (TCO). This is due to the use of common hardware for multiple network traffic types as well as direct cost savings from eliminating redundant management tools and systems for separate networks.

Past attempts at delivering a networking convergence technology did not materialize because the available solutions could not meet either the performance or investment protection requirements demanded by users. These failing technologies met one or the other of these requirements, but not both.

The adoption of 10-Gigabit Ethernet into mainstream data centers now makes it possible for Ethernet to match the raw performance characteristics of Fibre Channel (FC) Storage Area Networks (SANs). This development has given the industry the opportunity to develop a convergence technology based on the ubiquitous and familiar Ethernet networking standard that can also meet the needs of the storage networking market. This new technology is commonly referred to as Enhanced Ethernet (EE) or Converged Enhanced Ethernet (CEE). The standard that enables SANs to operate on Ethernet is known as Fibre Channel over Ethernet (FCoE). These emerging networking technologies will enable the next generation, lower cost, flexible, virtualized data centers of future.

Key Findings

The key requirements of a converged networking technology are:

- Sufficient aggregate performance to support all converged network traffic classes simultaneously.
- Seamless integration and protection of investments in existing networking hardware, software, and expertise.
- Enhanced Ethernet built on priority-based flow control and 802.1au Congestion Management will be the basis of the converged network.
- FCoE will provide the bridge between existing Ethernet infrastructure and FC-based SANs.
- The CEE network will increase Return On Investment (ROI) by enabling IT managers to cost effectively implement flexible systems that can meet the needs of high performance applications demanded by users.

Introduction

Data center IT managers have been burdened with managing multiple network fabrics, each of which addresses a different application. Ideally, IT managers prefer a simple networking technology that can be used for all types of network traffic, whether it is a Local Area Network (LAN), SAN, or Inter-Processor Communication (IPC). This white paper describes the trends driving the need for a "converged network technology", the benefits of such a technology, and the industry's answer to these needs.

The Case for Convergence

Early interest in network convergence was driven primarily by moving from proprietary hardware, software, and protocols to a standards based model that promised improved ROI based on the cost efficiencies of widely available hardware. Over time, IT infrastructure deployments showed a rapid growth and evolution in the number of specialized fabric installations, pushing them into the mainstream rather than niche application fabrics. With the characteristics of these "specialized" fabrics now virtually a requirement for all application deployments, it no longer makes sense to deploy multiple technologies. The driving need for converged fabrics is to reduce duplicate physical and logical networking resources in a growing data center environment. The following figure shows a typical non-converged data-center.

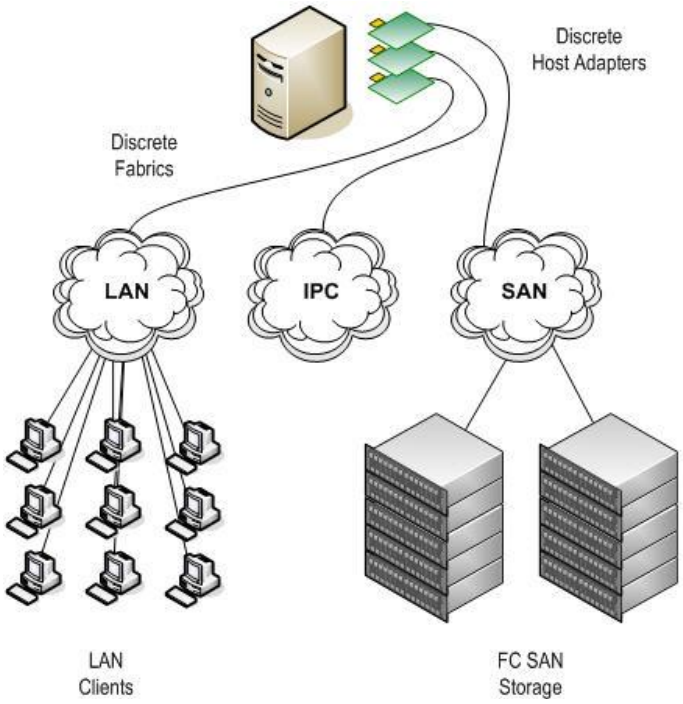

Figure 1. Shows a typical non-converged datacenter

Note that each server has at least three I/O adapters (six for full redundancy) and three discrete fabrics, all of which require their own management applications and personnel. This seemingly simple issue drives the TCO of a rapidly growing data center to unprecedented levels:

- The raw deployment cost is higher due to lower economies of scale for each fabric type.
- Operational costs are increased due to the need to manage raw electrical power, cooling requirements, and personnel for each separate network fabric.

A properly architected, designed, and deployed converged network will offer the following benefits:

- Reduced costs in hardware, physical resources, energy, and personnel.
- Increased flexibility to allocate network resources which, when further enhanced by server virtualization technologies, will enable increased ROI by ensuring that all resources are optimally utilized.

Barriers to Convergence

Since reliable enterprise data center operation is critically important to the success of large enterprises, data center IT managers are extremely conservative in their technology choices. As a result, unproven technology is slow to be adopted and is subject to long evaluation cycles. This is especially the case where systems interact and a problem in one area could critically impact another unrelated system.

Based on the evolution of LAN, SAN, and High Performance Computing (HPC) networks, coupled with the shortcomings of previous networks, a customer driven set of requirements for a converged network technology has emerged:

- **Raw Performance.** Each traffic class in a discrete network must have equivalent performance. For example, a minimum of 10Gbps is required to provide similar throughput as emerging 8Gbps Fibre Channel and 10Gbps Ethernet. Without the ability to provide comparable, if not better, performance levels than currently available, the technology's applicability to IT problems will severely limit the environments in which it can be used.
- **Investment Protection.** The converged network must allow the continued use of extensive investments in current network infrastructure.

  Investment protection includes the ability of the new infrastructure to provide a seamless migration path and interoperability in a mixed technology environment. Simply put, companies cannot afford to forklift upgrade to new systems or to have bottlenecks between old and new technologies.

  The converged network protocol must allow the continued use of investments in software, including hardware drivers and management platforms. Drivers must seamlessly integrate and interoperate with the converged protocol without needing to be upgraded. Both the investment in testing a

known driver and management software stack as well as the accumulated expertise of the IT staff must be preserved.

- **Management and Class of Service Isolation.** The converged network must be able to isolate each traffic class as if it was behaving as a discrete physical network, in terms of both performance and manageability. If there is any chance that one traffic class can affect another traffic class, reliability will be in question, thus rendering the system useless for business critical applications.

While the requirements for network convergence appear to be straightforward, previous attempts to achieve convergence through technologies such as InfiniBand and iSCSI have not materialized. The reasons that these technologies could not address the convergence requirements are outlined in the following paragraphs.

InfiniBand

- No investment protection. The advent of InfiniBand required a new infrastructure, including switches and complex bridges for interconnecting with Fibre Channel and Ethernet networks. In addition, it did little to preserve the SAN or LAN management layers. The resulting lack of investment protection and technology migration path has directed InfiniBand as a solution primarily for HPC applications.

1Gbps Ethernet/iSCSI

- No performance match. Predominant 1Gbps Ethernet was not comparable to the available 2Gbps/4Gbps Fibre Channel speeds or 10Gbps IPC network speeds that were available in the marketplace at the same time. IT managers could not use 1Gbps Ethernet when they already relied on 2Gbps or 4Gbps Fibre Channel speeds.
- No investment protection. At first, iSCSI seemed like an ideal protocol to consolidate storage on Ethernet. However, it was problematic due to the complexity of bridging iSCSI and FC SANs brought on by the additional overhead and implementation of TCP/IP offloads. In addition, iSCSI and Fibre Channel based management tools were not designed to work together. As a result, there was no leverage of Fibre Channel management expertise. With the re-use of Fibre Channel software and expertise impossible, there was no migration path for IT managers to take to iSCSI.
- No traffic class isolation. As currently defined, Ethernet has the ability to tag packets with Quality of Service (QoS) requirements and to implement flow control to prevent lost packets. However, these mechanisms are not designed to work at the same time to achieve traffic class isolation. As a result, a network can be either loss-less or implement QoS, but cannot be combined to implement a converged network technology that meets the traffic isolation needs of data center IT managers.

Building a Converged Data Center Network

With 10Gbps Ethernet and 8Gbps FC price and availability converging in the market, the industry has recognized that the time is right to develop a converged network technology that leverages these key data center networks. As a result, advancements and enhancements to Ethernet are now in development. These efforts, collectively referred to as EE or CEE (in the context of a converged network), are poised to deliver the promised benefits of a converged network technology that can overcome the prior barriers to convergence.

Performance

With 10Gbps Ethernet now a reality, a key shortcoming of Ethernet as a network platform for a convergence technology has been eliminated. 10Gbps Ethernet exceeds the performance of both emerging 8Gbps FC and multiple port 1Gbps Ethernet, making it possible for the latest Ethernet to replace any existing Ethernet or FC link on a throughput basis.

Investment Protection/Ease of Migration

To preserve existing investments in Fibre Channel, a method is required to seamlessly integrate Fibre Channel and Ethernet. To address this requirement, the new FCoE standard has been proposed by the industry.

Fibre Channel over Ethernet (FCoE)

FCoE is a new ANSI T11 standard for the encapsulation of a complete FC frame into an Ethernet frame. The resulting Ethernet frame is transported over Ethernet networks. The SCSI encapsulation in the FC frame does NOT require the implementation and processing overheads of TCP/IP based SCSI protocols like iSCSI and iFCP.

Fibre Channel is architected to provide reliable high performance operation in a loss-less network. However, Ethernet does not have such loss-less characteristics. One solution is to implement Fibre Channel on an Ethernet fabric that implements 802.3 Annex 31B pause-based flow control. This is a potential solution. However, it does not address the need for convergence of multiple traffic flows on Ethernet, as this flow control mechanism will stop ALL traffic rather than only the specific traffic that requires flow control. The industry is working on defining two standards to address these issues that, when combined, will make Ethernet a loss-less network with traffic isolation between classes of service.

The two additional technologies in development that will enable both high reliability and network convergence are Priority Based Flow Control and 802.1au Congestion Management. These technologies are described in the following paragraphs.

Priority Based Flow Control

Priority Based Flow Control enables separated traffic classes on a single fabric by enabling an Annex 31B pause "like" mechanism for each traffic class flow. This solution allows the designated higher priority traffic to flow at the same rate, while lower priority traffic may be temporarily paused when congestion occurs. As such, the differentiation in characteristics of traffic flows on an Ethernet network can occur, for example, to separate storage and LAN traffic or to separate SAN and IPC traffic.

IEEE 802.1au – Congestion Management

802.1au Congestion Management is the second key component required to enable a converged Ethernet fabric. In an Ethernet network, packet loss typically occurs due to network congestion. The network congestion is "solved" by dropping packets. Congestion Management will control packet insertion into the network to mitigate congestion in the fabric.

The Emerging Networking Marketplace

The development of a converged Ethernet network for the data center does not mean the end for discrete single purpose networks; each network technology will have a place in the deployment of IT services. In general, there will be three overlapping market segments in addition to a segment that will always require systems built on the highest performance discrete network technologies:

- The mainstream enterprise data center
- Small and Medium Enterprises (SME)/enterprise branch office
- Small and Medium Business (SMB)/branch offices

Each technology will deploy a combination of technologies to meet their needs. The following table illustrates a likely deployment scenario encompassing both discrete network technologies and converged 10Gb networks.

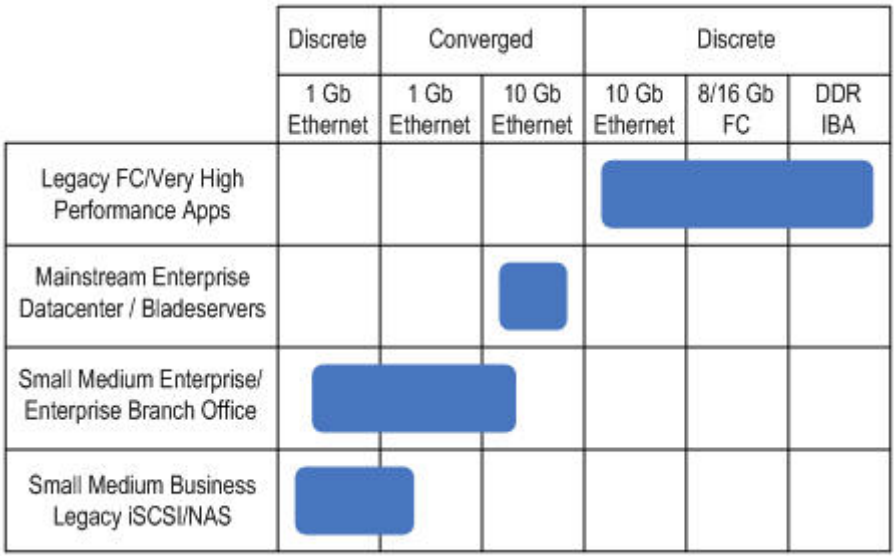

Figure 2. Deployment Scenario for Discrete and Converged Networks

Enterprise Data Centers

Enterprise data centers have the most to gain by deploying a converged Ethernet network due to the performance cost savings and flexibility offered by the converged network, especially when deployed in virtualized blade server environments. It is expected that a majority of new data center server deployments will implement a converged enhanced Ethernet.

Small Medium Enterprises/Enterprise Branch Offices

The IT needs of small and medium enterprises and the branch offices of enterprises are highly varied. Two common elements are smaller budgets and fewer IT staff to implement new technology, except where IT is a key part of a new business strategy. In addition, these organizations are not on the cutting edge of performance needs. For this reason, this segment will begin to leverage 10Gbps Ethernet later in time than the enterprise segments. When implementations begin to rollout, a mix of Fibre Channel and iSCSI SANs will be migrated to the CEE network for the high end of this segment, with smaller sites remaining on 1Gbps networks and taking advantage of the loss-less Ethernet features to improve performance.

Small Medium Business/Branch Offices

The SMB market has been slow to adopt Fibre Channel and has instead chosen Network Attached Storage (NAS) and iSCSI SANs based on 1Gbps Ethernet. It is expected that the performance and manageability of these smaller scale systems will continue to meet the needs of these users. As converged enhanced Ethernet switches move into the mainstream, these users will be able to leverage the loss-less characteristics to improve the performance of their applications and may choose to converge their networks, since iSCSI storage operates natively over Ethernet.

Conclusion

The emergence of 10Gbps Ethernet has spawned enhancements to a ubiquitous network that will enable true network convergence on a single hardware platform for multiple data traffic classes. This new network will enable data center IT managers to implement flexible systems that are more cost effective. In addition, these systems will meet the needs of high performance applications required by users and achieve a higher ROI. Each market segment will be able to take advantage of the new technology at a different time in its life cycle. Enterprise data centers will most enthusiastically embrace the technology to help simplify the installation and management of large server farms. Small and medium enterprises using Fibre Channel will follow when performance needs and management become an issue with the size of their deployments. Smaller enterprises and SMB segments will continue to use 1Gbps SAN and NAS on 1Gbps networks, but will gradually migrate to enhanced Ethernet and convergence as the technology is adopted and prevails.

Oracle is committed to developing practices and products that help protect the environment

Ethernet: The Drive for Convergence
March 2010
Author: Mike Wojnar

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com

SOFTWARE. HARDWARE. COMPLETE.