IDC
Analyze the Future

TECHNICAL BRIEF

Enabling Technology for Blade I/O Virtualization

Sponsored by: Hewlett-Packard

John Humphreys

December 2006

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

Introduction

In June 2006 Hewlett-Packard (HP) announced its next-generation blade chassis, the BladeSystem c-Class. The BladeSystem c-Class portfolio was designed to address some of the key total cost of ownership (TCO) issues facing today's datacenter, including server management costs, utilization, and power and cooling. As part of the BladeSystem c-Class, HP introduced three new technologies:

- HP Virtual Connect Architecture
- HP Insight Control Management
- HP Thermal Logic Technology

These technologies play a central role in reducing overall datacenter operating expenses. They also differentiate the BladeSystem c-Class system both from competitive blade offerings and from rack-optimized servers.

In this technology brief, IDC examines HP Virtual Connect and the benefits and challenges associated with using this technology to virtualize I/O with HP BladeSystems. It joins the IDC white paper *Forecasting Total Cost of Ownership for Initial Deployments of Server Blades*, which describes the overall TCO advantages of the BladeSystem c-Class portfolio, as well as the technology briefs *Enabling Technologies for Blade Management* and *Enabling Technologies for Power and Cooling,* which focus on the operational and cost advantages of HP Insight Control and HP Thermal Logic, respectively.

Server and Connectivity Growth Contributes to Growing IT Management Costs

The explosion in servers in today's datacenter has led to exponential increases in data-transfer volumes and bandwidth bottlenecks. This growth poses significant problems for IT managers as they struggle to identify more cost-effective, scalable ways to manage the increasing complexity of servers and their associated LAN and SAN environments.

Over the past 15 years, two primary factors have driven the dramatic increase in the worldwide server installed base:

- **Expanding server footprint.** Today, most core business process applications are supported by multiple servers. It is not uncommon for applications to be supported by two, three, or even five or more servers per application. Further, with new processes such as email and Web-based applications constantly coming online, these additions only increase the number of servers an organization must support.

- **Shift in the server mix.** Twenty years ago an organization might have purchased a handful of mainframes and could have reasonably expected those systems to handle all of its IT needs. Customer buying patterns have evolved, however, as pressure built to reduce initial acquisition costs and new lower-priced technologies continued to be introduced. With each technology transition, the cost of acquisition came down by an order of magnitude, and today, x86 systems average about $4,000 per system. These lower price points reinforce the proliferation of applications, which can be deployed on their own dedicated servers. As a result, it is not uncommon for datacenters to support 5,000 or more servers, most of which support a single application.

Figure 1 demonstrates the effect of this installed base increase on IT budgets. Worldwide spending on new server acquisition has actually declined over the past 10 years; after peaking at $69.5 billion in 2000, it was $55.4 billion in 2005. In contrast, spending on server management and administration has risen from $22.7 billion in 1996 to $102.8 billion in 2005, and IDC projects it will continue to grow. As a result, management and administration costs now greatly outpace the cost of server acquisition.

FIGURE 1

Worldwide Server Installed Base, Acquisition Cost, and Management Cost

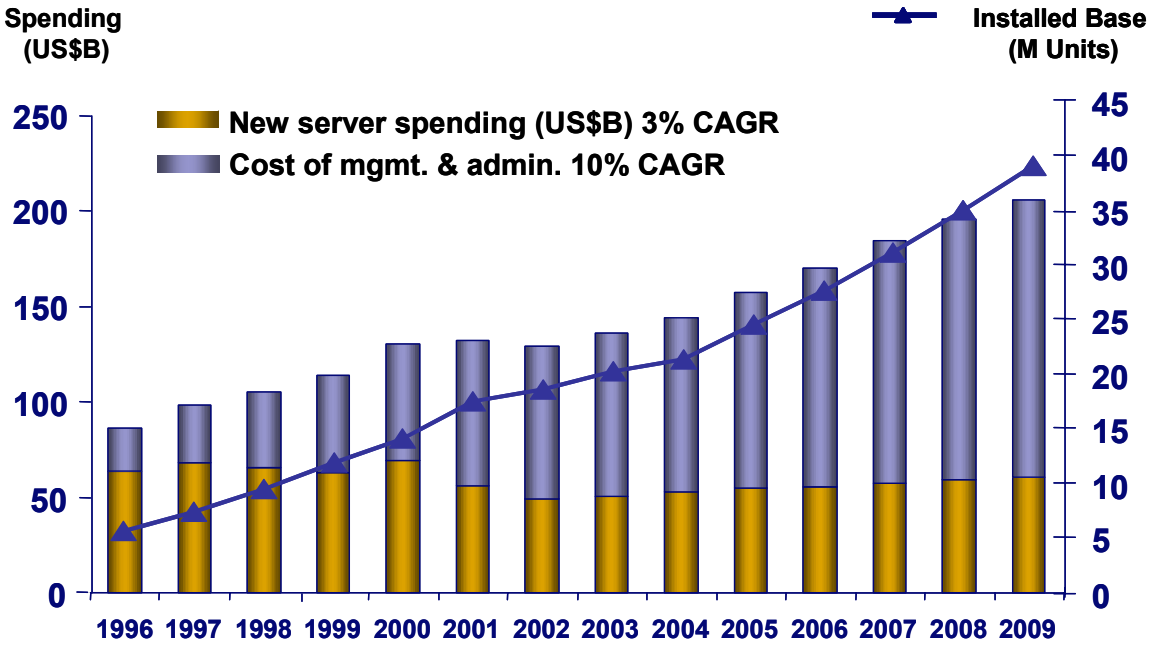

Source: IDC, 2006

Server sprawl poses additional challenges to IT organizations, including:

- **Inflexibility.** As currently deployed by most organizations, server configurations are static, hardwired, and difficult to change.
- **Manual coordination.** Too many people and too many manual steps are required to change applications or server configuration.
- **Underutilization.** With 1:1 application-to-server deployment ratios, only a fraction of most servers' total capacity is typically utilized.
- **Overprovisioning.** Along with underutilization, overprovisioning results in wasted rack space, along with increased power, cooling, bandwidth, and operational costs.
- **Unique management requirements.** Many servers support unique processes and applications, requiring IT organizations to maintain individual tools and coordinate inconsistent configurations to handle them all.

The bottom line is that server infrastructures consume large amounts of space, power, and IT administration time. The costs associated with deploying, monitoring, and managing servers have escalated to the point where IT organizations are actively seeking tools to help them manage costs by streamlining management and automating key server management tasks throughout the datacenter.

Current Approaches to Server I/O Are Costly and Inefficient

One of the key drivers of server management costs is the interdependency between servers and datacenter networks, including LANs and SANs. Previous-generation approaches inhibit efficient utilization of the network fabric as connectivity remains attached to the server or, in the case of blades, the chassis backplane. As a result, network administrators must be engaged every time a new blade is installed or an existing blade is replaced, which requires a lot of scheduling and wastes the time of all IT departments involved.

Furthermore, previous-generation blade servers suffer from Fibre Channel rates that limit usefulness in working environments. To accommodate peak traffic flow, IT administrators have to overprovision connectivity, resulting in wasteful, underutilized network infrastructures.

This situation results in higher-than-necessary TCO, due to both the waste associated with overprovisioning and the personnel costs required to closely manage all components of these complex fabric infrastructures.

HP BladeSystem c-Class

HP has made providing solutions to these issues central to the design of its BladeSystem c-Class. Specifically, HP has taken steps to deliver tools to help customers address today's operational pain points while putting themselves on a path to creating their own next-generation datacenters (see Figure 2).

FIGURE 2

BladeSystem Innovations

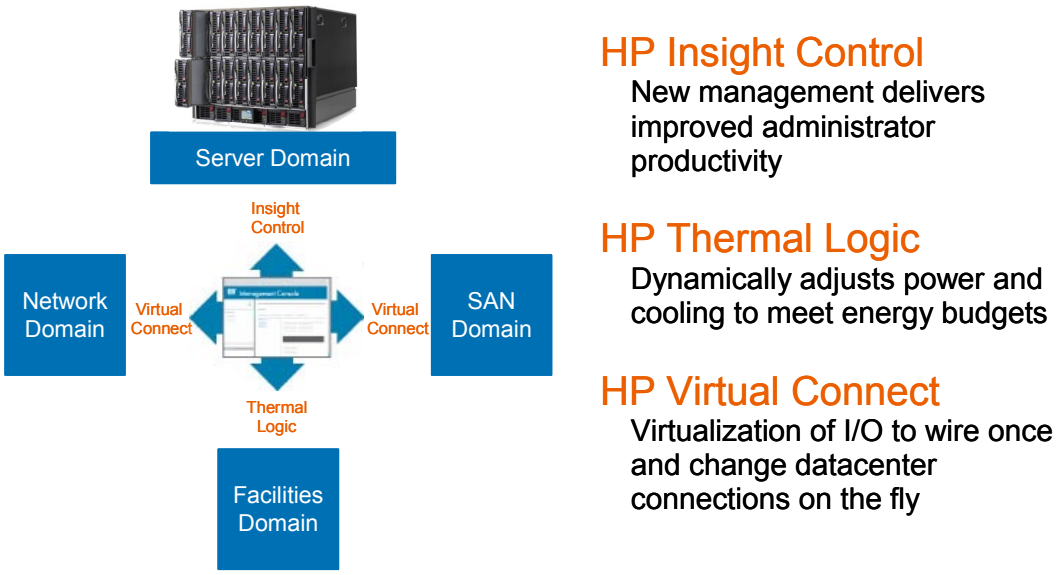

Source: HP, 2006

Core design principles include modularity and choice — enabling customers to adopt key features and integrate them when they are ready. HP is focusing on systems management with HP Insight Control, virtualization with HP Virtual Connect, and power and cooling with HP Thermal Logic.

This approach enables:

- **Just-in-time provisioning.** BladeSystem c-Class servers are preprovisioned and wired once. When reconfiguration is required, it can be done quickly and easily.
- **Automated coordination.** Domains and people are isolated from the upheavals of change.
- **Virtualization.** Devices and connections are managed as pools of resources, shielding the administrator from the complexity of managing individual servers.
- **Lights-out, "1:n" management.** Group management capabilities streamline processes and reduce management overhead.
- **Capital management and efficiency.** Indications are that BladeSystem c-Class servers are less expensive to buy and own than conventional IT server infrastructures.

HP Virtual Connect

HP Virtual Connect technology decouples servers and the network so that changes in the server infrastructure don't require changes in the network. It does this by aggregating connectivity into a single resource pool that can be shared physically or virtually across c-Class blades. Viewing connectivity resources in this way allows customers to simplify network connections, avoid switch management, and increase efficiency in operational management. Moreover, consolidating connectivity enables organizations to increase overall bandwidth utilization.

Virtualizing server I/O can help IT managers address the escalating costs of managing a large pool of servers in their organizations and minimize the necessary "touch points" for server and network administrators to add, move, and replace servers. Users can quickly and easily move resources to meet changing needs in their infrastructures, such as stopping bottlenecks or scaling to accommodate a spike in demand.

How Virtual Connect Works

Virtual Connect virtualizes BladeSystem connections to LAN and SAN environments. It defines a server connection profile for each server bay in an enclosure before the server is installed. This profile establishes the media access control (MAC) addresses for all network interface controllers (NICs) and the World Wide Names (WWNs) for all host bus adapters (HBAs). The SAN boot parameters then hold the MAC addresses and WWNs constant in software so that they remain unchanged, even if the bare-metal server is changed. Those same MAC addresses and WWNs are also assigned to a new server that replaces an existing server. Profiles can also be moved from one bay to another bay, facilitating rapid substitution of a new physical server for one that has failed, without disturbing any MAC address or WWN mappings. The LAN and SAN never see any changes. This decoupling of servers from the network provides pragmatic benefits to end users. For example, if a server is changed or a new server is added, the LAN and SAN already know how to communicate with it. Moving a profile to a different bay involves a mouse click. Changing the network connections on a profile also involves a few mouse clicks. So failing over to a spare server, moving a server from a development environment to a test environment (with different LAN and SAN connections), or moving a group of servers to a different application subnetwork requires just a few mouse clicks by the server administrator (see Figures 3 and 4).

FIGURE 3

Server Migration and Failover: Before

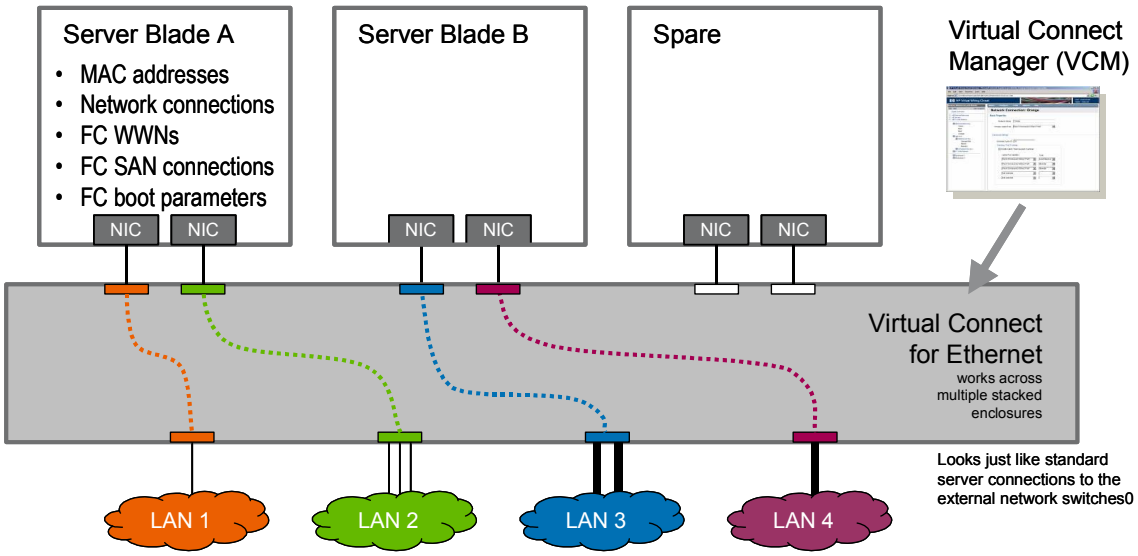

Source: HP, 2006

FIGURE 4

Server Migration and Failover: After

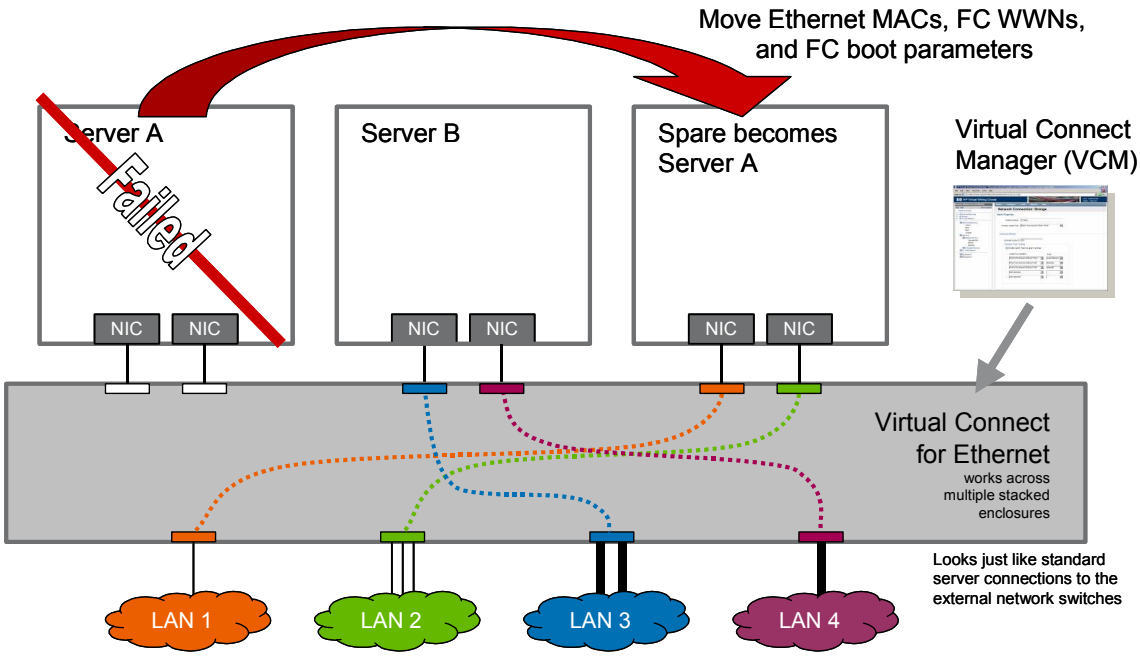

Source: HP, 2006

Virtual Connect greatly simplifies the management of server network connections, but more importantly, after deployment, the server administrator can quickly accomplish all of the previously mentioned functions without the help of network or storage administrators. Server, network, and storage administrators can be more productive, and it is not necessary to hold up the server provisioning process while waiting to schedule network support staff. Customers deploying Virtual Connect can run up to four simultaneous redundant fabrics with an aggregated bandwidth of 5Tb sec^{-1} or 160Gb sec^{-1} per server, providing scalability. Further, support for 10Gb Ethernet, 1Gb Ethernet, and 4Gb Fibre Channel is already built in.

Ethernet and Fibre Channel connection modules will plug directly into the BladeSystem c-Class interconnect bays to reduce cables without adding switches to manage. Pass-through modules can also be used to provide 1:1 server-to-network connections.

Benefits

The benefits of a virtualized approach to server I/O in the BladeSystem environment include:

- **Reduces server provisioning and management effort.** The virtualization of I/O helps reduce the time spent on issues associated with moving, adding, or changing servers and the associated costs. With Virtual Connect, the network can be preprovisioned and the network identity lives with the slot in the chassis, not the server. Because Virtual Connect saves each server profile and the server is stateless from a network perspective, organizations can get back up and running with a new bare-metal server very quickly, without loading drivers and so forth.
- **Simplifies networks.** The number of Ethernet, Fibre Channel, and management cables; Ethernet aggregation layer switch ports; and Fibre Channel director-class switch ports and transceivers — and their costs — can be reduced by more than 90%. Depending on the specific network configurations, cable reductions of 75% or more are possible. In addition, once the initial chassis configuration is complete, the time and effort required to deploy or configure new and existing servers are greatly reduced. Further, cable reduction can be achieved without adding any switches to manage.
- **Simplifies server connections.** Virtualizing I/O allows changes to the servers without affecting the datacenter view. SAN and LAN administrators have to provision the network only once and no longer need to help with server moves, changes, or additions. This virtualized approach gives them back the time they would have spent on these tasks and frees up server administrators from delaying changes while coordinating schedules with other groups.

Virtual Connect may be used to streamline efficiency in a variety of usage scenarios, such as test and development environments where servers need to be constantly scaled, configured, and reconfigured; datacenter scaling to address new capacity requirements; and business continuity and failover by streamlining the process of setting up remote datacenters and failover sites.

Finally, using Virtual Connect to incorporate additional I/O and more efficiently use the available I/O addresses a key barrier to customers that want to leverage blade systems across all datacenter tiers.

IDC Analysis

The HP BladeSystem c-Class enables IT managers to improve overall TCO by reducing the amount of effort required for server management and administration, including not only the time required to provision a new server but also the time required to provision network connections. Designed with an understanding of how complex IT infrastructures are actually managed, Virtual Connect dramatically reduces the time and effort required to provision and reconfigure server connections with storage and Ethernet networks.

Ultimately, Virtual Connect reduces TCO as a result of a number of factors. Virtual Connect allows IT managers to reduce the number of administrators required to manage LAN/SAN connectivity, increase flexibility, reduce time spent managing switch infrastructure, remove expensive travel costs associated with LAN/SAN management, increase switch fabric resilience through redundancy, and remove costly error margins by removing human intervention.

Key barriers to adoption exist, however. Some are at the BladeSystem level, including the need to prove the general applicability of blades across IT organizational needs and mission-critical applications. The market is just reaching the beginning of the mainstream adopter phase, and IDC believes that HP, along with other blade vendors in the x86 market, must overcome the market perception of blades as the exclusive domain of specialty solutions for large infrastructures.

Other issues at the I/O virtualization level include the need for customers to choose appropriate networking technologies, as well as the perception that Virtual Connect will introduce new complexity, best practices, or networking standards into customers' datacenters. In fact, Virtual Connect is a straightforward use of industry-standard Ethernet and Fibre Channel that looks like a collection of NICs to Ethernet networks and like HBAs to Fibre Channel fabrics. The challenge for HP will be to communicate how this new technology actually reduces complexity and simplifies the job of network and server administrators rather than adding to their management workload.

Foundation for Next-Generation Virtualized Datacenter

Virtual Connect is a foundational technology for the next-generation datacenter. IT organizations will increasingly embrace virtualization in their infrastructures, allowing creation and management of compute resource pools that can be easily provisioned and changed to meet organizations' fluctuating demands and service-level commitments. Virtualization of all IT resources will play a role, including server, storage, and I/O virtualization, and implementing Virtual Connect today provides one of the building blocks for the fully virtualized infrastructure of tomorrow.

Choosing Appropriate Networking Technology

Historically, the integrated networking of a blade solution has been positioned as a benefit to customers. Increasingly, IDC is finding that the inclusion of switch and server in the same chassis is causing conflict in that network and server administrators are challenged over ownership issues. Virtual Connect can help to resolve both ownership and switch management challenges. The opportunity for HP is to offer choice with respect to networking options. With the c-Class product, customers can now choose from a pass-through, a switch, or the Virtual Connect option for either Ethernet or Fibre Channel networking.

Concerns About Introducing Additional Complexity into the Network

Another barrier that must be overcome is that of the perceived lack of standards in blade servers and concerns over any additional complexity that customers may be introducing into their networks. IDC research shows that standards are important to mainstream technology users because standardization ensures price competition, reduces choice in multivendor environments, and safeguards interoperability. On the pricing side, the dynamics of the blade marketplace are already addressing the drive for commoditization. The incorporation of Virtual Connect is a clear point of network differentiation for HP and has the potential to offer significant value to customers that have complex or rapidly changing server environments; however, many customers may perceive that Virtual Connect acts as a switch and may be concerned that introducing it into their networks will compromise their best practices. The challenge for HP is to illustrate to customers that Virtual Connect does not cause disruptions or outages and that it preserves complete compatibility with networking standards and best practices currently adopted by organizations.

Customers May Require Prove-Out

Virtual Connect represents a new approach to virtualizing server I/O, and as a result, customers are likely to approach the technology cautiously. To alleviate this concern, HP will need to clearly illustrate the value of the technology and showcase best practices gathered through actual customer implementations.

Conclusion

As the next generation of HP blade server architecture and technology, the BladeSystem c-Class provides a significant advancement in reducing overall server management complexity and TCO. One of the technologies enabling these efficiencies is Virtual Connect, which virtualizes I/O connections to SAN and LAN environments, enabling organizations to wire once and add or replace servers or change datacenter connections on the fly.

By incorporating HP Virtual Connect into future infrastructure deployments, customers can improve TCO by reducing the time required to deploy, manage, and maintain their server deployments, eliminating the need to involve SAN and LAN administrators when adding, reconfiguring, or redeploying servers. Further, they can increase bandwidth efficiency by enabling management of I/O resources as a single logical pool while reducing physical cabling requirements and contributing to more efficient server cooling.

Copyright Notice