

## WHITE PAPER

---

# Rearchitecting the Datacenter Infrastructure for Improved Performance and Efficiency

Sponsored by: IBM

---

John Humphreys

Lucinda Borovick

April 2004

## EXECUTIVE SUMMARY

Growing competition and pressures to reduce costs and improve profitability have put a renewed focus on the performance of enterprise datacenters. To maintain business-critical operations and provide their organizations with a competitive edge, enterprise CIOs are seeking to create a simpler and easier-to-manage datacenter topology with an infrastructure that increases application availability and performance, reduces support costs, improves security protection, and adjusts quickly to unpredictable workloads and changing business needs.

The typical datacenter environment is riddled with complexity, due to the broad inventory of equipment and the countless number of relationships among the various system and network components built up over the years. CIOs who want to streamline operations and provide more secure and reliable services may be uncertain of where to start in unraveling the complexity of years of IT build out.

As a first step in rearchitecting their infrastructure for greater simplicity, efficiency, and flexibility, many organizations are consolidating their rack-optimized datacenter servers into a blade server chassis, and then integrating the storage fabric and Layer 2 switches into the chassis. Blade servers were originally viewed as a means of conserving datacenter space and power, but they are now valued more for their ability to increase availability and manageability and to lower acquisition and ownership costs. Integrating Layer 2 switches into the chassis further enhances the consolidation and manageability benefits of a blade server, and provides the necessary backplane to support many server workloads.

Moving beyond datacenter recentralization and physical server consolidation, it is becoming increasingly clear that CIOs will need to consider not just the server or hardware platform, but the entire IT environment when attempting to improve flexibility, enhance resiliency, increase uptime, or otherwise improve service-level metrics for the business services that IT supports. Additionally, enterprises must protect their datacenters from disruption and put in place an on-demand infrastructure foundation to improve operations for renewed business growth. Datacenter managers need infrastructures that can reduce support costs, allow them to improve service levels, integrate and deploy applications faster, and respond quickly to changing business needs and market conditions.

To that end, IBM has teamed with Cisco to jointly leverage their expertise and lay out a solution path that satisfies both the short-term tactical and longer-term strategic

needs of enterprise customers. Early efforts have resulted in the combination of IBM's BladeCenter with Cisco's Layer 2 embedded switch as a platform for consolidation. More strategically, the two companies are also working to integrate their business service initiatives to create a single-reference architecture that an enterprise can leverage for IT optimization or to improve service delivery.

## CIO CHALLENGES IN THE DATACENTER

Running all the facets of a complex business requires a vast amount of IT resources. To manage these resources effectively, CIOs often have to consolidate them, either through the centralization of datacenters and/or the physical fusing of hardware, data, and applications. Additionally, CIOs must strive to make IT more cost-effective through better utilization of hardware, software, network, and staff, while accommodating fluctuating business conditions that call for maximum flexibility in IT resources.

Users' demands on the IT organization are also increasing. Users want an IT infrastructure that is highly available, secure, and dynamically scalable to meet the ever-changing operational demands on the business. At the same time, users increasingly demand delivery of their IT resources with a variable rather than a fixed cost, enabling them to readily compensate for the business "variability" they regularly encounter.

IDC believes the challenges facing CIOs have already begun to shape the evolution of datacenters and associated technologies, leading to more virtual and dynamic computing environments that maximize provisioning and management automation, and are simple and cost effective to deploy and manage.

Tangible practices that have proven helpful to CIOs and IT managers in their efforts to realize this vision include consolidation, integration, process automation, datacenter virtualization, and network transformation.

- ☒ **Consolidation.** The net effect of consolidation — especially with blade servers — is a flattening of the datacenter topology by eliminating layers of infrastructure (both servers and network switches). As a result, the infrastructure is easier to manage because the topology itself is less complex.
- ☒ **Integration.** Assimilation of disparate technologies results in a higher-performance computing environment that is more comprehensive, but simpler and easier to manage.
- ☒ **Process automation.** To become more productive, IT administrators need to offload their mundane, day-to-day tasks to software and hardware solutions that can "automate" certain tasks, allowing them to focus on higher-value projects, such as planning for new application and service rollouts. Process automation offers further productivity improvements. By offloading processes to the automation software, administrators are able to accomplish tasks more efficiently and deliver changes in near real time.

- ☒ **Datacenter virtualization.** Beyond the provisioning of operating systems, patches, and applications, IT administrators are increasingly being tasked with developing a strategy that enables the company to continue to add capacity to the IT environment for little or no additional cost. This challenge is being addressed by virtualization or grid-computing offerings, which enable enterprises to drive up hardware utilization, effectively adding capacity to the system without having to purchase new hardware by better harnessing what is already there.
- ☒ **Bandwidth upgrades.** The first path to increased application availability is via network bandwidth upgrades. To satisfy the increasing demand for bandwidth in the datacenter, and throughout the enterprise, companies are upgrading their network to Gigabit Ethernet. In some organizations, 10Gigabit is already deployed to build larger and simpler connections in the datacenter, replacing multiple gigabit interconnects.
- ☒ **Network convergence.** Converged networks enable enterprises to take advantage of technology advances, standards, and increased price/performance curves in ethernet deployments, and to leverage staff expertise in handling new types of traffic, such as voice and storage over the ethernet network. In time, convergence will be the foundation for new applications.

While this picture of a technology utility is a long-term, strategic vision, IDC believes that, by focusing on the few core principles of automation, efficiency, security, and accountability, CIOs and IT administrators will be able to successfully navigate the crowded and uncertain waters of an IT environment in transition. These core principles define a predictable path that leads toward more integrated, on-demand architectures. IDC further believes that blade servers are a critical step in the progression toward that vision.

## EVOLUTION OF THE DATACENTER

In the past, server and networking companies generally worked independently to address CIO needs in the datacenter. However, CIOs increasingly view their datacenters as a tightly coupled infrastructure that provides critical business services, creating a need for server and networking companies to work more closely to deliver integrated platform solutions.

Blade servers became popular in datacenters as a way to lower power consumption, as well as means to reduce floor space and rack space requirements when both space and power were the constraining issues. Today, concerns about space and power are less critical than the need for a consolidated platform that fits into a scale-out strategy to minimize acquisition and operational costs while providing higher levels of availability, performance, scalability, security, and manageability at the service level.

Today's blade products are designed to deliver dramatically improved levels of manageability, automation, and virtualization through a combination of embedded, customizable hardware and software components, such as network switches, management cards, and provisioning solutions. Additionally, IDC believes the combination of blade servers with management and automation solutions can

significantly lower infrastructure acquisition costs and the ongoing operational expenses.

---

## **How Will the IBM and Cisco Alliance Address Datacenter Requirements?**

CIOs who have deployed blade technology cite a number of benefits. Some mention the ease of use of the products, while others stress the serviceability or the modularity of the product. Still others like the increased density, which allows them to fit more computing capacity into their already crowded datacenters. One message common to the CIO experiences is that blades bring the benefits of consolidation.

Embedding and integrating a LAN switch within the blade chassis adds to these consolidation benefits. With this combination, CIOs are not only consolidating many servers into one physical box, they are also consolidating a layer of networking. The entire package becomes a single managed entity, which increases reliability and reduces the operational costs associated with managing multiple servers.

In keeping with customer need and demand for such integrated offerings, IBM and Cisco have integrated the Cisco Systems Intelligent Gigabit Ethernet Switch Module with the IBM BladeCenter (see Figure 1). This effort embeds and integrates Cisco intelligent Ethernet switching directly into the BladeCenter chassis, facilitating enhanced server and network virtualization. This helps IT organizations to immediately realize the benefit of the collaborative datacenter architecture within the chassis itself.

In developing its BladeCenter solution, IBM focused on a number of specific attributes that stem from the consolidation benefits that blades offer. These attributes include:

- ☒ **Application availability and resiliency.** In an increasingly complex datacenter environment, customers are demanding ways to deal with changes while maintaining high availability. Hot-swappable blades, power supplies, and fans that include inherent traffic redirect available from an integrated switch platform provide the first step in assuring high availability.
- ☒ **Consolidation.** While many enterprises have consolidated their server and storage architectures, they have yet to address the need for consolidation in their network architecture. The integrated switch is a start to bringing consolidation benefits to the network architecture.
- ☒ **Network leadership.** Cisco has significant presence in enterprise networks. By integrating a Cisco switch within BladeCenter, IBM customers can extend their Cisco networks into the blade chassis. By peacefully coexisting within a Cisco network, the integrated solution makes deployment easier, and it allows the BladeCenter to take advantage of Cisco services that are deployed throughout the network.
- ☒ **Manageability.** Streamlining the infrastructure management process improves the operational efficiency of the datacenter and reduces the demand on the IT administrator. IBM has recognized that by consolidating servers and switches

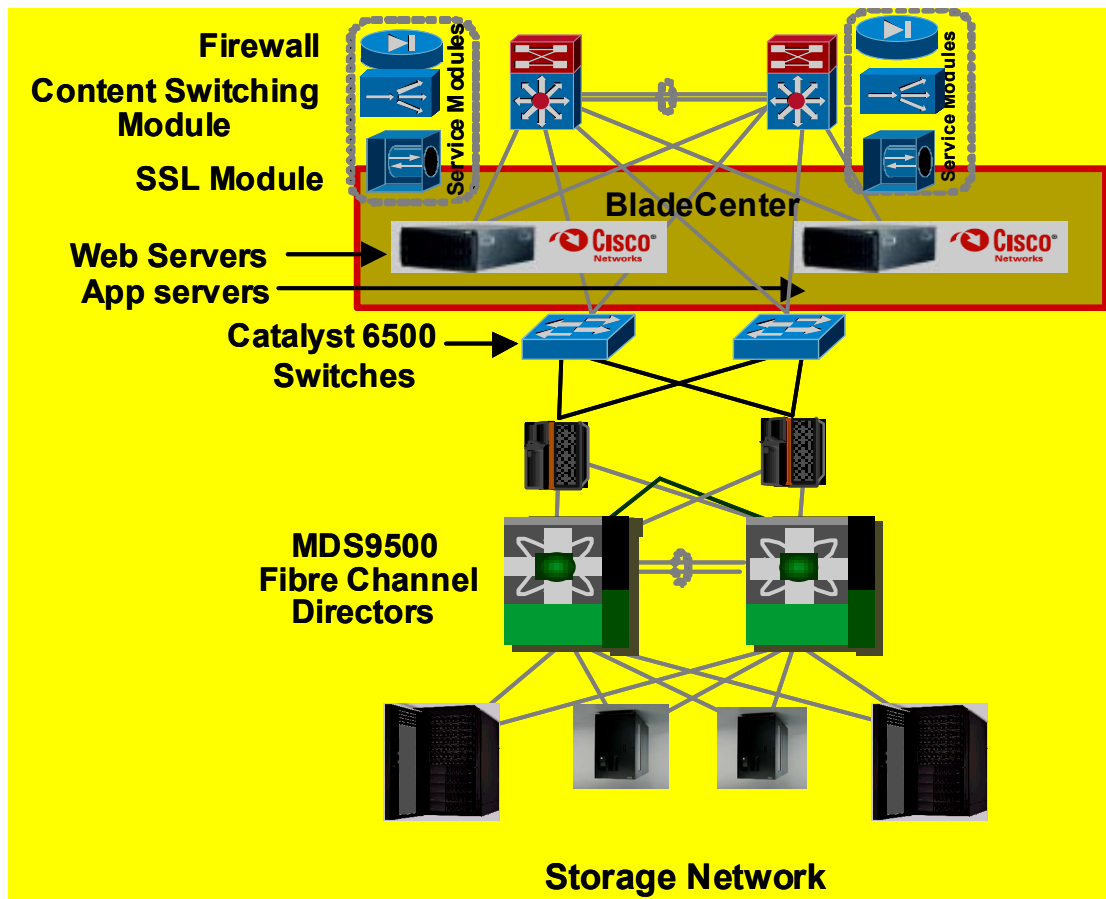
into a chassis. Blades allow customers to centralize and physically consolidate servers and switches to lower IT operational and management costs. Providing full interoperability into existing Cisco and IBM datacenter products further enhances manageability. Customers have flexibility in managing the integrated switches with options that include BladeCenter management interface, IBM Director, a Web browser or Cisco command line interface (CLI). The switch can be managed by IBM Director software, IBM's flagship product for improving datacenter manageability. IBM Director comes standard with the BladeCenter. It allows users to manage the entire device holistically, while also delivering granular monitoring and management of the individual server, switch, or component on the device. IBM includes integrated systems management processors on all the blades and on the chassis-based redundant management module, providing aggregated management and advanced features, such as streamlined blade software deployments. Other features that can improve overall manageability include interoperability with SNMP network management platforms, like CiscoWorks, and other Cisco offerings that aid in inventory, debugging, and troubleshooting the network.

- ☒ **Customization.** The ability to right-size blade performance by matching blade/processor to a specific workload is helpful. IBM delivers extra versatility by allowing customers to mix and match not only blades within the chassis but also switches, management modules, power supplies, and cooling units to configure a highly customized platform for meeting particular needs.
- ☒ **Security.** The integrated gigabit Ethernet switch module includes network components that block unauthorized users from gaining network access. Features include centralized access control, port security, and user authentication.
- ☒ **Reduced cost of acquisition.** Blades offer a pay-as-you-go model of scaling resources. Adding blades is less expensive than adding rackmounted servers, since components such as switches, power supplies, cooling fans, and hardware interconnects are shared across all the blades. Because server blades share resources, the acquisition cost compared with individual standalone devices is generally less once the chassis is roughly half full. In the case of IBM BladeCenter, if the organization has a need for more than six or seven servers, a blade solution can reduce upfront costs, while still providing headroom for future growth.
- ☒ **Reduced cost of ownership.** The most significant opportunity to contain and lower the cost of business services at the end user's site exists in the effective management of a large datacenter. IDC estimates that, on average, approximately 60–70% of server life-cycle costs are associated with its management. By offering enhanced application availability, more comprehensive ecosystem manageability, easier scalability, and a lower cost of acquisition, IBM aims to offer a platform that can deliver a significant reduction in the total cost of ownership (TCO) relative to standalone, disparate, and patchwork approaches to building an IT environment.

In interviews with IT executives at enterprises that have deployed blade servers within their datacenters, IDC found that the combination of server and LAN switch consolidation delivers more benefits than just server consolidation or LAN switch consolidation alone. Enterprises in the IDC study that migrated from rack-optimized servers in the datacenter to blade server configurations will be able to reduce their average cost of ownership for server and network infrastructure by 48% over three years through server consolidation and reduced server acquisition and maintenance costs. Embedding Layer 2 switch functionality within the blade server chassis would reduce these costs by an additional 33%, for a total savings of 65%, by bringing similar consolidation and maintenance savings to the networking infrastructure and by further reducing acquisition costs.

**FIGURE 1**

Datacenter Consolidation: Where the BladeCenter Collaboration Fits



Source: IBM, 2004

## THE FUTURE OF THE DATACENTER

As discussed earlier in the paper, CIOs and IT managers are looking beyond point solutions and increasingly viewing the datacenter more holistically to address challenges on the business service level. Broadly, these concerns include the following:

- ☒ **Increased datacenter scalability.** This scalability starts within the blade server chassis and between multiple blade server chassis, but rapidly expands to encompass a virtual pool of servers in the datacenters. These servers may be tied together in a grid to handle peak workloads and increase response times.
- ☒ **Better utilization of resources.** In improving the cost of ownership, customers are looking to squeeze more from their existing infrastructure. One way to do this is to pool resources and break down application silos by virtualizing both the underlying resources and overlying applications.
- ☒ **Improved application availability.** Enhancing application uptime through intelligent provisioning of resources increases critical business service levels.
- ☒ **Enhanced service levels.** Tying demand and supply information for a given application or business service to the priority that the application or service has to business users empowers IT managers with better decision-making capabilities, and allows them to leverage the information to improve not only the efficiency of the business, but the profitability as well.
- ☒ **More granular accounting.** Increasing the measure and control functions within the datacenter provides a better understanding of the day-to-day, month-to-month, or year-to-year variations in the business, as well as a unique look into the consumption of services, and how and whom to charge for supplying those services.
- ☒ **Greater security protection.** Security is an overarching problem that spans every inch of the network. As the types of corporate network users expand to include partners, suppliers, and an increasingly mobile, wireless-enabled workforce, security concerns are heightened, and the need for application-level security becomes a top priority.

It is through these lenses that IDC believes enterprises will be making decisions on whom to partner with to provide their datacenter solutions in the coming years. Enterprises are increasingly valuing the relationship with their enterprise technology provider and looking to these partners for guidance through the complexity that is today's IT environment.

---

### Addressing Broader Datacenter Requirements

Beyond just the tactical integration of the BladeCenter with a Cisco Layer 2 embedded switch, the two companies have also announced the formation of an alliance in charge of tackling the challenges and concerns of the datacenter administrator from a holistic perspective. The two firms recognize that building an on-demand datacenter will necessitate evolutionary changes in the network. Specifically,

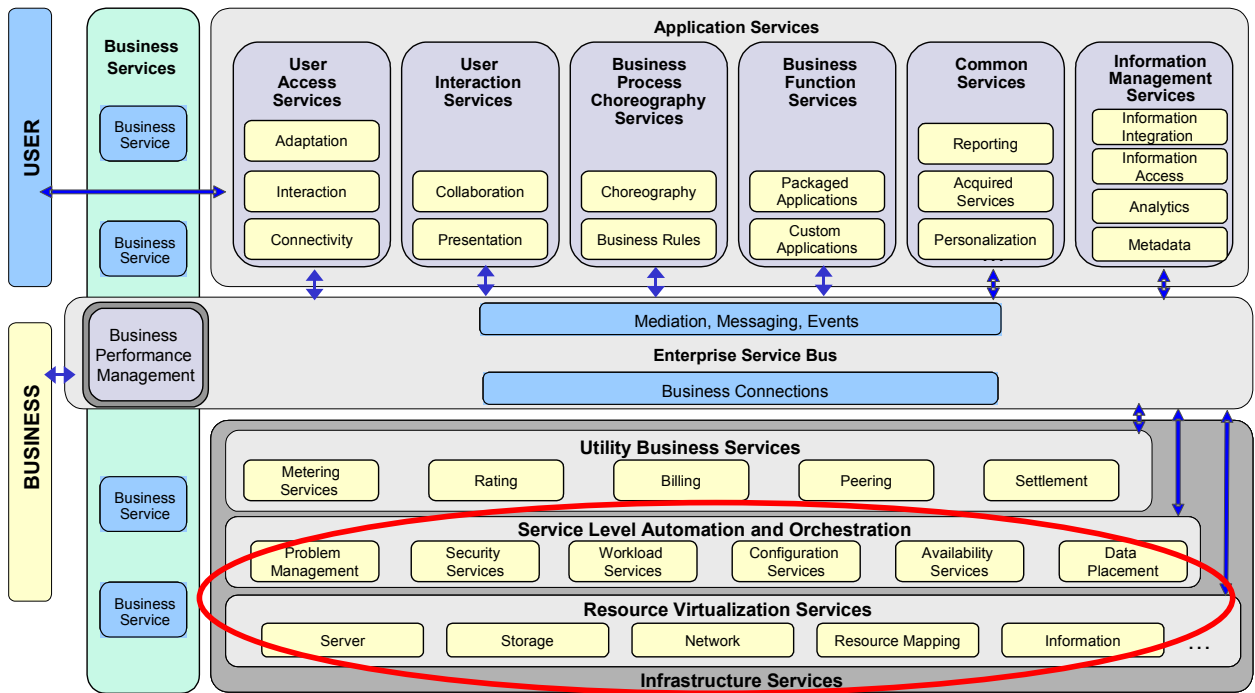
IBM and Cisco are working to integrate their respective On Demand and Intelligent Information Network initiatives to deliver a reference architecture that offers customers a means to address their broader datacenter requirements.

Functionally, the alliance between IBM and Cisco is designed and charged with integrating the Cisco Business Ready Data Center with the IBM On Demand Operating Environment. The optimization of these two independent initiatives is focused on the key principles of business transformation, providing technology innovation, enhancing customer productivity, and enabling new business models. The initial area of focus is on providing infrastructure solutions — specifically solutions around the areas of resource virtualization, datacenter orchestration and service-level automation.

Figure 2 illustrates where in the On Demand Operating Environment the IBM and Cisco collaboration is currently focused.

**FIGURE 2**

Initial Areas of IBM-Cisco Collaboration



**Initial area of IBM-Cisco collaboration**

Source: IBM, 2004

IBM and Cisco are squarely set on helping customers increase the efficiency of their datacenters through a combination of better utilization of the hardware in the datacenter (resource virtualization), automation of mundane manual administrator



tasks (datacenter orchestration), and increased application availability (service-level automation). The partnership has also identified security as a principle tenet — in that the solution must provide these efficiency benefits without compromising security.

### ***Resource Virtualization***

Leveraging resource virtualization, IBM and Cisco are evolving the datacenter architecture through a combination of existing, new, and emerging products, technologies, and services. The primary focus of the IBM and Cisco relationship is IT simplification, including an infrastructure that is easier to provision and manage, as well as the creation of a single, consolidated, logical access to all the available resources in a network. Current collaborative initiatives in the resource virtualization area include the integration of the IBM eWLM and Virtualization Engine with the Cisco Business Ready Data Center Network.

### ***Datacenter Orchestration***

After the resources and associated services in the datacenter are virtualized. Cisco and IBM can turn their attention to datacenter orchestration. Specifically, the companies are working to add systems-provisioning workflows for Cisco Business Ready Data Center Network components within the Tivoli Intelligent Orchestrator and Provisioning Manager. This effort will yield provisioning capabilities that extend beyond just a server or network switch to ones that can be applied more on a business service level.

### ***Service-Level Automation***

The On Demand Operating Environment uses autonomic computing to enable the infrastructure with a sense-and-respond characteristic, so that the customer can tie the underlying infrastructure to the overlying business priorities. In this way, IBM and Cisco can not only match supply of services with the demand for those services, but they can also help to document business priorities so that the most critical services are given the highest level of availability. This is accomplished by treating each application and resource as a service. These services have well documented interfaces that have been defined according to industry standards, permitting both the services and underlying resources to be dynamically provisioned to where they are needed most and where they will be most effective.

### ***Security and Availability***

Critical considerations associated with consolidating resources onto a pooled infrastructure are security and availability. Cisco's in-depth integrated security, high availability, and business-continuity networking technologies provide a resilient foundation capable of helping to protect application, server, and storage resources, and the network itself. This resilient datacenter network foundation addresses availability requirements within a single datacenter, as well as among multiple datacenters, promoting higher levels of business continuity and extending the On Demand Operating Environment capability to encompass all datacenters within the enterprise network.

## **CHALLENGES/OPPORTUNITIES**

The challenges facing IBM and other server vendors as they attempt to transition their product lines to map to customer demands include developing a complete product portfolio, defining a succinct value proposition for the pragmatic market majority, continuing to drive down the cost of computing and networking, and more broadly expanding customer perceptions on their place inside the datacenter. Key to all these challenges is the development of a future roadmap for the customer that is both pragmatic in the near term and encompassing in the longer term.

As blade server vendors endeavor to penetrate beyond early adopters and technology enthusiasts, it is clear that demand will flow toward modular designs that do not forgo network intelligence, computing power, or availability over density. Dual-processor and even quad-processor blade designs are highly desired, and as such they have an opportunity to become central to tomorrow's datacenter infrastructure.

The integrated IBM and Cisco On Demand Operating Environment endeavors to bring a new level of security protection, resiliency, performance, scalability, availability, on-demand capability, simplified management, and TCO benefits to not just blade server deployments, but to the entire datacenter. The challenge for the IBM and Cisco alliance is to continue to strike a balance between tactical offerings and strategic vision to remain relevant as it goes forward. The vision must include a portfolio of new application-specific network services that include security, quality of service, and service-level priorities. The alliance has begun to address these challenges by articulating a strategy that emphasizes a coexistence with the existing networking infrastructure in the network.

## **CONCLUSION**

IDC believes that increased consolidation provides enterprises with the ability to unravel the complexity of their datacenters, and that the combination of server and LAN switch consolidation provides more benefits than just server consolidation or LAN switch consolidation alone. There is also a trend for technology customers to adopt a scale-out computing model that leverages a pay-as-you-grow model of hardware acquisition, which leads us to conclude that the blade server platform is well positioned as a major platform for the next generation of computing. Blades have already established momentum with early adopter customers and are beginning to become more widely deployed. Beyond the ease of use, ease of deployments, and density benefits, current users of the technology see blades as a platform that delivers computing and network consolidation, which translates into both reduced acquisition costs and a lower TCO. The IBM BladeCenter offering is now customizable with an embedded and integrated Layer 2 Gigabit Ethernet switch module designed by Cisco.

Besides the TCO savings of the platform, IBM and Cisco have announced joint development of the On Demand Operating Environment. The intent of this program is to integrate Cisco's Intelligent Information Network with IBM's eBusiness On Demand initiative to deliver a set of integrated technologies that span the compute and communication networks and deliver improved application availability and resiliency, better scalability, and improved security protection across the entire datacenter

environment. The joint development is critical because building an on-demand architecture will necessitate changes in the network. As the leader in enterprise networking, Cisco is uniquely positioned to provide strategic guidance to customers looking to migrate to utility computing.

Early efforts have been successful in coupling IBM's eWLM with the Cisco Content Services Module, in creating system provisioning workflows for Cisco's Business Ready Data Center Network within the Tivoli Intelligent Orchestrator, and in developing Tivoli SAN Manager enhancements to better manage the Cisco MDS9000 within the IBM SAN Volume Controller. These efforts seem to indicate not only the depth of the IBM and Cisco alliance, but also the strong focus by the two companies on working to create differentiated value for their customers.

---

### **Copyright Notice**

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2004 IDC. Reproduction without written permission is completely forbidden.