



## WHITE PAPER

# BLADE SERVERS IN THE CISCO BUSINESS READY DATA CENTER—BRINGING CISCO NETWORK INTELLIGENCE INTO EMERGING SERVER ARCHITECTURES

**As the enterprise has become more dependent on networked applications and resources to increase productivity, the data center has emerged as the most business-critical place in the network. The trend of data center consolidation to improve service levels and responsiveness to business needs has led many enterprises to both centralize and virtualize their computing resources, moving away from monolithic servers and toward computing and storage resource components that can be flexibly consolidated and pooled by a single network resource. This provides the opportunity to reduce operational costs and expenditures, and motivates data center administrators to take an architectural approach to their data centers to provide a secure and resilient IP infrastructure that best optimizes applications and data resources.**

## UNDERSTANDING BLADE SERVERS

As computing resources become a shared pool rather than a dedicated resource, blade servers have emerged as a cost-effective way of consolidating management, configuration, and power capacity into one platform. A blade server is a chassis that houses many servers, or computing resources. Typically, blade servers are used to replace older server farms where density is a requirement, or where new applications that use clustering are being deployed. This form-factor change—combining many CPUs into one platform—allows the data center manager to reduce operational costs and save valuable rack space. Blade servers are a relatively small portion of the server market today, encompassing less than four percent of the server units and one percent of the server revenue. According to IDC, this is expected to grow by 2007 to 27 percent of server units and up to 13 percent of the server revenue. Many of the leading server vendors, including IBM, HP, Sun Microsystems, and Dell, currently have blade servers on the market.

Blade servers address issues related to server consolidation and space management. The consolidation of many servers into one enclosure can provide significant gains in rack space in the data center. Cable management becomes considerably simpler—fewer cables are needed to connect into the blade server chassis. However, these benefits often come with some tradeoffs. Blade servers are still relatively new in the market and have not yet taken *all* environmental considerations into account. For example, many servers in a single enclosure require a greater amount of cooling. This could limit the number of blade servers in a particular area or rack, as too much warm air can circulate around the systems. Also, blade server enclosures, particularly in a large data center, have been found to strain the limits of the data center's raised flooring. This forces data center managers to disperse the blade servers, eliminating the advantage of rack space savings.

Blade servers offer the integration of switching technology into the chassis, which provides local switching between servers and uplinks to the rest of the data center. When you are integrating a blade server and switch module, you need to consider the existing data center network topology. Although the switches are a form-factor change only, the aggregation or distribution of network services, insertion of Layer 3 technologies, and subnet implementations are important considerations in maintaining high availability and stability. Planning and building the architecture around present and future intelligent switching capabilities, such as those found in Cisco® Catalyst® switches, is important to the deployment of a scalable, secure, and resilient data center that will provide investment protection over time.

Cisco Systems® supports customer choice, and will help customers as they consider network management, optimal vantage point for selected network services, stability, and performance. An architectural approach to the data center network enables optimization of networking services for both standalone server and blade server environments. Cisco is working closely with some blade server

providers to test and recommend deployment optimizations within the context of the Business Ready Data Center architecture, and will continue to evolve these partnerships over time.

## **INTEGRATING BLADE SERVERS INTO THE BUSINESS READY DATA CENTER**

The Cisco Business Ready Data Center brings an architectural approach and end-to-end intelligence to the data center by integrating security, resilience, and application optimization into each layer of the network. This holistic approach enables the entire data center, including blade servers, to deliver service to users, applications, and communication systems and creates a high-performance solution to better enable customer business applications.

There are numerous considerations for the deployment of blade servers into a networked data center architecture, including integration of security into the IP infrastructure, end-to-end resilience, placement of Layer 4 to 7 services, and manageability and reliability. These capabilities must exist as part of the network architecture, whether the network is independent or integrated into the blade server.

### **Integrated Security**

The data center is an integral part of the business, and security is critical. Security integrated into the network must cover three distinct areas—secure connectivity, threat defense, and trust and identity management. Blade servers, which exist in the access layer of the data center architecture, have unique requirements for security that must be taken into account. This is typically limited to establishing secure access to the switch, protecting against attacks, and identifying users and servers attempting to access the network. Many capabilities, such as firewalling, are consolidated in the aggregation switch for easy configuration, management, and control.

First, access to the switch itself, whether embedded into the blade server or external, must be secured. This requires specific features, such as Secure Shell Protocol Version 2 (SSHv2) or RADIUS. Both protocols are means of securing connectivity to access the management interface of the switch. Without these capabilities, it is considerably easier for an unauthorized user to change the configuration, hindering the blade server's ability to deliver service to its constituents.

The access layer of the data center must also defend against potential attacks. The switch must provide capabilities by which to prevent, as well as defend against, network attacks. Cisco Catalyst Intelligent Switching provides unique capabilities to help ensure that attacks are detected and stopped. First and foremost, the switch must allow the data center manager to configure access control lists (ACLs). These ACLs provide the first layer of defense by governing who has access to the server (in this case, blade server) resources. The next layer of defense prevents unauthorized connectivity and snooping of data. This is provided in Cisco Catalyst switches using features such as:

- *Port security*—Helps to ensure that a limited number of MAC addresses are learned per port
- *MAC address notification*—Informs the network manager if a MAC address has suddenly moved (as would happen if an end station were being spoofed)
- *Dynamic Address Resolution Protocol (ARP) inspection*—Ties a port to an ARP request, helping to ensure that a default gateway cannot be spoofed
- *IP Source Guard*—Protects against an IP address being spoofed

Finally, the access layer switch must make sure that the network only allows authorized devices to connect to it. This is handled by Cisco Catalyst trust and identity management capabilities. Using the IEEE 802.1x authentication protocol, only devices that correctly authenticate with a RADIUS or TACACS+ server are allowed connectivity to the network. This helps to ensure that servers are authorized on the network. Future capabilities will tie 802.1x authentication to other security mechanisms, such as being sure that an end station is using the latest virus software, thereby limiting the data center's exposure to spreading viruses. In all, Cisco Catalyst Integrated Security helps to ensure that the servers, either blade server or standalone, are safe from both internal and external attacks.

## Availability and Resiliency

The availability of the data center infrastructure, as well as its ability to recover from a fault, help to ensure that the network resources and applications will be “on” 24 hours a day, seven days a week. This requires the network to provide two levels of resilience—device-level resilience and network-level resilience. Device-level resilience refers to the ability of the device to prevent or recover from a failure. There are several features in the Cisco Catalyst switch that help ensure device availability. Integration of a switch in the blade server does not allow for some of the most basic device-level resilience capabilities, such as redundant control processors; this capability requires a separate, standalone access switch. Other features help ensure the availability of the switch platform, such as CPU rate limiting, which is available on the Cisco Catalyst 6500 Series. This feature limits the amount of traffic that can be sent to the processor of the switch CPU, which is critical in the event of a denial-of-service attack.

Network-level resilience refers to the capabilities of the network to recover from a device, link, or protocol failure while maintaining the availability of the network. This involves both protocol considerations and network design considerations. Many blade servers can integrate dual switching engines and then “dual-home” the servers within the chassis. This integration can potentially complicate resilience mechanisms, such as the Spanning Tree Protocol.

Using “pass-through” mode on the integrated switches, which facilitates bypassing the internal switches and connecting the network interface card (NIC) or host-bus adapters (HBAs) directly to external switches, can diminish this complication. However, the switch must provide many of the enhancements Cisco has created innovations for spanning tree, including UplinkFast and the standard version, 802.1w. This provides subsecond failover, either within the blade server or in the access layer. The integration of Layer 3 switching into the network in the distribution and core provides interaction with the rest of the enterprise network, providing reachability and best-path forwarding. Layer 3 switching is not enabled in the blade server and access layer—that brings the failure domain of the core deep into the data center. Network availability is an architectural consideration, not a device-only consideration. The optional integration of switching into the blade server introduces concerns that can be mitigated by intelligent switching.

## Delivery Optimization

Different applications have different requirements. The network must be able to provide the capabilities that these applications require, either to operate efficiently or to better map the enterprise business priorities. Cisco Catalyst Intelligent Switching brings to the data center architecture three technologies that optimize the delivery of applications—quality of service (QoS), multicast, and content switching. These end-to-end technologies each have different components that map to different locations in the network.

Blade servers provide the access into the network for a given number of servers. The application data, as soon as it enters the network, needs the network to perform certain functions to optimize its delivery. The first is classification of data according to application requirements or business priorities. QoS at the access layer requires that this classification take place, in hardware, and then map into the appropriate queue so that QoS can be applied. Queue scheduling using weighted round robin and strict priority helps ensure that applications get more (or less) bandwidth on the network as required.

Many applications—particularly video, but also many financial applications—use multicast for efficient distribution of data from one source to multiple destinations. At the access layer, Internet Group Management Protocol (IGMP) snooping receives and processes IGMP Join requests issued by the application so that, at Layer 2, multicast traffic is efficiently transported to the next-hop router or Layer 3 switch. By implementing this function close to the server (or within the blade server), bandwidth conservation can be achieved.

Another important component of optimizing delivery of applications is content load balancing. Load balancing allows for better use of server resources and processing power by balancing connections from the network to the server farm. This capability is not a blade server requirement, but an architectural requirement. To best use *all* the server resources, the data center architecture typically places the load-balancing services in the aggregation switch, such as a Cisco Catalyst 6500 switch. In most cases, adding content load-balancing services in the blade server adds complexity while not using all the server resources in the data center.

## **Enhanced Manageability**

As with any other place in the network, management is an important consideration. The switching infrastructure must provide the appropriate hooks for the network management platforms to allow the network manager to configure, monitor, and troubleshoot the network. This, in turn, requires the devices in the network to be “manageable” so that the appropriate information can be relayed between the network manager and the network device.

Blade servers pose no new manageability challenges—the chassis would be managed as would any other server platform. However, if a switching device is embedded into the blade server, it must then be managed as the rest of the network infrastructure would be. In theory, this should not be an issue; however, few blade server switches are built with manageability in mind. They are typically integrated simply for server connectivity.

Cisco brings several manageability features into intelligent switching. The first, most common capability is a common command-line interface. This reduces the operational cost of managing different user interfaces and eases the configuration of the platform. Cisco Discovery Protocol is integrated into all Cisco devices, enabling Cisco device manager software as well as any Cisco connected device to know what Cisco device it is connected to and to get basic information about it (such as its IP address). This aids in both monitoring and in troubleshooting basic connectivity. For more advanced troubleshooting, port mirroring or Switched Port Analyzer (SPAN) is provided to allow a network manager to “span” a potentially troublesome port to a management port to which a Remote Monitoring probe might be attached. Remote SPAN allows a data center manager connected to one switch to “span” traffic on another switch.

Manageability and management are systemwide considerations. Cisco Catalyst Intelligent Switching helps to ensure that blade server ports and connected switches can be managed and maintained effectively, as would any other Cisco networking device.

## **CISCO CATALYST INTELLIGENT SWITCHING AND BLADE SERVERS**

The need to integrate intelligent switching capabilities into the blade server is important for the deployment of a systems-level architecture. Cisco is working with partners to integrate Cisco Catalyst Intelligent Switching technology into the blade servers, thereby providing the right hooks in the blade server to enable a scalable systems approach to data center design and deployment. Whatever a customer chooses for server deployment Cisco will provide solutions for the best network infrastructure. Cisco is committed to refining networking infrastructures to deliver the next generation of data center products and to increase IT productivity.

## **SUMMARY**

Blade servers are not built for networking. As a result, current offerings do not provide the intelligent switching capabilities required of a switching platform. These include capabilities such as advanced QoS, integrated security, and network availability. The Cisco Business Ready Data Center brings intelligence to the networking infrastructure, enabling the deployment of a scalable architecture. Integration of switching technology into the blade server must provide the right hooks to enable a scalable systems approach, which in turn delivers high performance and service to the network. Cisco is committed to working with its partners to integrate its switching innovations and use its experience to enhance the data center architecture, thereby combining the benefits of blade servers with superior switching technology.

**Corporate Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
www.cisco.com  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 526-4100

**European Headquarters**

Cisco Systems International BV  
Haarlerbergpark  
Haarlerbergweg 13-19  
1101 CH Amsterdam  
The Netherlands  
www-europe.cisco.com  
Tel: 31 0 20 357 1000  
Fax: 31 0 20 357 1100

**Americas Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
www.cisco.com  
Tel: 408 526-7660  
Fax: 408 527-0883

**Asia Pacific Headquarters**

Cisco Systems, Inc.  
168 Robinson Road  
#28-01 Capital Tower  
Singapore 068912  
www.cisco.com  
Tel: +65 6317 7777  
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on **the Cisco Web site at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica  
Croatia • Cyprus • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR  
Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico  
The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia  
Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan  
Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2004 Cisco Systems, Inc. All rights reserved. Cisco, Cisco Systems, the Cisco Systems logo, and Catalyst are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0403R) BG/LW6347 0504