



Blade Server Benefits

Evaluating new server technology in an IT environment

Results from an Intel IT proof-of-concept (PoC) project show that blade servers offer significant cost savings in acquisition, operations, and facilities while presenting no major barriers to deployment.

May 2004

Executive Summary

Faced with shrinking budgets due to the current economic downturn, Intel IT needed to lower costs without reducing our overall level and quality of services. We chose a radical new approach that would cut costs and increase management efficiencies, while still enabling us to respond quickly to increasing and changing business needs for computing, networking, and storage capabilities.

Intel IT needed to lower costs without reducing our overall level and quality of services. We chose a radical new approach...

This paper describes our proof-of-concept (PoC) study to explore the benefits and risks of integrating blade server technology into our existing infrastructure. We also looked at physical installation and power utilization of blade servers. For this effort, we recruited a company-wide team of experts in areas such as storage area network (SAN), networking, systems management, platform, and data center technologies to carry out the study.

Our study identified no significant barriers to introducing blade servers into our environment. At the same time, we *found significant benefits to deploying blade servers*, including:

- Lower acquisition costs
- Lower operational costs for deployment and troubleshooting and repair
- Lower power, cooling, and space requirements
- Dramatically reduced cabling requirements
- More efficient out-of-band management

A total cost of ownership (TCO) analysis for our environment showed potential reductions of:

- 10 to 20 percent for acquisition costs
- 3.8 to 10.6 percent for deployment costs and troubleshooting/repair
- 8.5 to 25.3 percent for facilities costs

Based on our PoC study, we are now preparing to deploy blade servers throughout the Intel IT computing environment.

Contents

Background	3
Why Blade Servers?	3
What are Blade Servers?	3
Blade PoC Methodology	4
Challenges and Unknowns	4
PoC Tests and Results	4
Installation Testing	5
Power Utilization Testing	5
Integration Testing	5
Network Functionality	6
SAN Interoperability	7
Clustering and Failover	7
Blade Chassis and Out-of-Band Management	7
Application Execution	8
Summary of PoC Results	8
Business Benefits	8
Deployment Guidelines	10
Conclusion	11
Acronyms	11

Background

Why Blade Servers?

The provisioning capability of our existing rack-mounted server environment is both lengthy and inflexible. Our search for a better approach led us to blade servers, an emerging technology that promised reduced costs in areas such as capital expenditures, operational expenses, and physical plant requirements while improving efficiency. Blade servers allow for up-front provisioning of the chassis and switch components when adding compute blades as needed. This provides a more dynamic provisioning model, which results in just-in-time provisioning as well as the ability to work around data center freeze periods

What are Blade Servers?

Blade servers form the basis for a modular computing paradigm. A blade is a single-board server that contains one to four processors, memory, local disk storage, and on-blade network interface card (NIC) and SAN connectivity. Figure 1 on the next page depicts the insertion of a blade into a blade chassis.

As shown in Figure 2 on the next page, a blade chassis may also hold one or more Ethernet and SAN switch modules, one to four power supplies, one to two shared management modules, and cooling resources. Chassis components communicate across a fully redundant midplane, enabling hot swap functionality of the chassis components and easy serviceability.

Figure 1. Blade server chassis (front)





Figure 2. Blade server chassis (rear)



Blade PoC Methodology

The goal of the study was to collect information on which to base recommendations for how blade servers should be positioned and deployed at Intel. We wanted to determine the technical benefits, limitations, and risks of blade server technology and determine what changes would be required to integrate blade servers into our environment.

Our PoC team included representatives from our IT global operations and engineering groups, the IT architecture group, and the finance department. We also brought in end-users from applications engineering to participate on the team.

The team included experts in SAN, networking, systems management, platform, and data center technologies, who developed and executed test plans for their respective areas of expertise. The test system was set up in a single location in the U.S. Members of the team, who were located throughout the world, used remote access to complete their parts of the test plan.

Challenges and Unknowns

During the planning phase, we identified several key questions:

- **What effects would deploying blade servers have on our data center architecture?**

We wondered if the data center would require any retrofits to accommodate blade servers, such as changes to power or cooling systems, standard server racks or structural modifications to the data center floor to support changes

in density and weight. We also wanted to determine how cabling aggregation would impact the data center.

- **What are the key integration issues?**

We needed to explore how to integrate blade servers into our existing Ethernet network, Fibre Channel SAN, and systems management structure.

- **How would integrating blade servers into our infrastructure impact reliability and stability?**

We knew we would need to leave much of the testing of reliability and stability for the engineering phase. The specific product and configuration selected for final deployment would have a significant impact on the overall reliability and stability of the deployed system, and product selection was beyond the scope of the PoC.

However, reliability and stability are critical issues, so we had to take special note of any issues that might affect availability.

PoC Tests and Results

Our PoC team created a test plan and assigned tasks to the different test teams. Each test team accessed the system remotely to conduct their assigned tasks. The teams compiled their test results in a central location.

The plan included procedures to test installation, power utilization, network functionality, SAN functionality, clustering and failover capabilities, out-of-band server management, and application execution. Specific tests and results are described in the sections that follow.

Installation Testing

Installation testing focused primarily on physical installation. Operating system provisioning, for example, was outside the scope of our PoC and was not tested.

We physically installed and cabled 14 blade servers, network, storage, and management modules, keyboard, video, and mouse (KVM), and power in a blade server chassis. We then loaded a standardized build of a commercially available server operating system on each blade server and configured the network access.

We also completed a hot removal and installation of a blade server and swapped two blade servers between slots, making sure, in each case, that the reinstalled blades booted properly.

All installation tests completed successfully. The most noteworthy finding was the significant reduction in the time to install the blade environment versus the rack environment, and with greater functionality. We completed the physical installation and chassis cabling for the 14 blade servers in 2.5 hours—including external network and storage switches—compared to 8.5 hours to install and cable 14 rack-mounted servers without external network and storage switches. Installing these switches would add approximately 1.5 hours to the physical provisioning of the rack server.

Cable aggregation contributed substantially to this improved efficiency, as shown in Table 1. Our blade server test system required 21 cables to interconnect Ethernet and SAN switch modules, power, and KVM. A rack-mounted server system with equivalent functionality would have required 154 cables, a 733 percent difference.

Power Utilization Testing

For the power utilization tests, we reconfigured the 14 one-processor blades in our test system as seven two-processor blades. We then compared the power utilization of our blade server test system to that of a rack-mounted server with a similar configuration. We used a commercially available benchmarking tool to generate a central processing unit (CPU) load for each blade.

Measured at the 20 percent and 90 percent CPU utilization levels, power consumption test results showed significantly improved power utilization of blade servers compared to rack-mounted servers (see Figure 3 on the next page). For example, at 90 percent CPU utilization, the blade server system containing

Table 1. Number of cables required

Cable Type	Rack-mounted Server System	Blade Server System
Ethernet	56 ¹	10 ²
SAN	28 ³	4 ⁴
Power	28	4
KVM	42	3
Total Cables	154	21

1. Four Ethernet cables for each rack-mounted server: one for applications, one for backup/recovery, one for clustering, and one for out-of-band management
2. Only ten Ethernet cables for the blade server chassis, including two additional cables for redundant management modules, due to aggregation of functionality
3. Two SAN cables for each rack-mounted server for redundancy
4. Only four SAN cables for the blade server chassis, due to aggregation

seven blades used 1248 watts while an equivalent system of rack-mounted servers would have used 3422 watts (extrapolated from one to seven two-processor rack-mounted servers).


We tested CPU utilization with one, three, five, and seven dual-processor blade units, which allowed us to extrapolate power utilization out to 14 blades. We tested power utilization for one rack-mounted server and extrapolated that data to determine the power utilization for 14 rack-mounted servers. The wattage-consumed lines depicted in Figure 3 on the next page are based on the duration of our test cycle, which was approximately 15 minutes (the dashed portions indicate extrapolated data).

Integration Testing

One of our primary concerns, and a major focus of our testing, centered on how blade servers would integrate into our existing environment. For example, what impact would integrating blade servers have on our Ethernet and SAN infrastructures, and how would deploying blade servers affect our data center power and cooling requirements?

The constraints of our lab environment somewhat affected the range of tests we could include in our PoC. More complete functional and performance testing incorporating failover capability for the chassis components will be performed in the engineering phase of the project prior to production deployment.

Figure 3. Power consumption versus number of dual-processor units



Network Functionality

To understand the impact of introducing blade Ethernet switches into our existing infrastructure from a functional performance and operational perspective, we completed testing in the following areas.

Switch Interoperability

- **Link-level testing.** Connected the blade switch to the layer 2 Ethernet switch and tested the link for rich operability.
- **Auto-negotiation.** Tested auto-negotiation at speeds of 10 Mbps, 100 Mbps, and 1,000 Mbps between the blade switch and the existing infrastructure switch.
- **802.1q trunking.** Tested trunking of the virtual local area network (VLAN) between the blade switch and the existing network infrastructure.
- **Spanning tree configuration.** Tested the spanning tree process on the blade switch.

All performance and functional tests passed. However, auto-negotiation between the blade center switch and our existing network infrastructure didn't occur automatically. We had to set the speed and duplex through the NIC driver. We also noted that if we introduce a different supplier's model switch into our environment, we would need to update our operational procedures.

Blade Server Switch Performance

We tested throughput performance by downloading test TCP (TTCP) on all server blades and testing the performance of the blade server switch. Throughput performance was 510 Mbps, which was well within the expected performance range.

Management Functions

- Simple network management protocol (SNMP) function of the blade server switch
- Reporting of port traffic statistics
- Trap and syslog configuration
- Port mirroring
- History log
- Remote monitoring of the blade switch through HTTP/HTTPS and Telnet
- Upgrade and downgrade of the blade switch software
- Configuration file manipulation (restoring an old configuration file)
- User account maintenance

All tests passed except the configuration file manipulation test. Configuration files were in binary format, which made manipulation and restoration of a configuration file difficult.

We also found that the blade switch we used had no command line interface (CLI). The lack of a CLI is an operations issue we would have to work around for this particular switch, since en masse scripting changes in a large enterprise environment such as Intel's require a CLI.

Failover Functionality

We successfully tested rebooting of a single Ethernet switch in the chassis and determined the convergence time.

SAN Interoperability

We conducted tests in the following areas to assess SAN interoperability and ability to boot the operating system from the SAN.

Integration of Blade and Existing SAN Switches

We had to set the configuration of our existing SAN switches to interoperability mode to communicate with the integrated blade chassis SAN switch.

We tested the effect of an interoperability mode change on zoning to ensure that information was retained. Configuration information for zoning was maintained.

Blade Server Connectivity to Storage

Our tests showed that the host bus adapter (HBA) driver loaded properly, zoning was functional, and the allocation of storage was functional.

Booting from SAN

We tested that the SAN configured as a boot device, that we could install the operating system on the SAN drive, and that the server booted properly.

Although the procedure for configuring the HBA as a boot device was complicated and not well documented, the configuration was successful. It required a change to the boot-up order and modifications to the HBA BIOS.

We did not observe any improvement in boot time with the operating system installed on the SAN drive.

Lab limitations prevented us from testing dynamic multi-pathing. We will complete this testing in the engineering phase of deployment.

Clustering and Failover

Clustering for failover is an important part of our environment, so we clustered a blade server with an available rack-mounted server through a public network interface. We would not typically cluster dissimilar hardware architectures, but these were the test resources that were available for the PoC.

Once installed and configured on the blade server and the rack-mounted server, the cluster services successfully failed back and forth between the two servers.

Our testing was not extensive enough to develop a recommendation for a cluster configuration for our environment. We will complete this step in the pre-deployment engineering phase.

Blade Chassis and Out-of-Band Management

We conducted tests in two system management areas: blade chassis management and out-of-band management (remote control).

Blade Chassis Monitoring and Management

- **System status.** Tested that the chassis and components, including individual blades, Ethernet switches, management module, power modules, blowers, and temperature display the correct status.
- **Event log.** Tested that all critical and warning events were logged and time stamped.
- **Alert management**, including:
 - SNMP messages at the management module – Verified that event error messages were passed to the management module.
 - SNMP trap management – Tested the ability to send traps to the SNMP host and e-mail.
- **LEDs.** Tested the capability of the management module to toggle LEDs on and off at the front panel.
- **Hardware/firmware vital product data (VPD).** Tested that vital data about all the blade system hardware could be viewed, edited, and captured.
- **Power/restart.** Tested that each blade server could be powered on and off individually and that all blade servers could be powered on and off together.
- **Ethernet switches.** Tested restart of the Ethernet switches from the management module. Set up and configured switches using a web interface.

All the above features worked as expected.

An added benefit was the ability to control, view, and configure the status of components in the blade chassis from the shared view provided by the blade chassis management module.

Working from a shared view was much more efficient than logging on to each individual server, as is required in our rack-mounted server environment.

Out-of-Band Management (Remote Control)

We tested the following remote-control features:

- Shutdown and restart
- KVM control
- Firmware and BIOS updates on groups of blades
- Power on and configuration of system BIOS for the blade
- Configuration of boot device sequence for individual blades and all servers

Some performance and reliabilities issues, such as dropped connections and user interface problems, occurred during out-of-band management over the wide area network (WAN). Also, the management module did not authenticate with the directory service we were using. Everything else tested successfully.

The ability to perform out-of-band management on multiple blades from a shared view provides an efficiency benefit compared to rack-mounted servers. In addition, the aggregation of management cards—one for each blade chassis rather than one for each rack-mounted server—and associated cabling has the potential to yield significant savings in acquisition costs.

Due to the performance and reliability issues found during PoC testing, we recommend using out-of-band management facilities only when the operating system is down and using in-band facilities otherwise. We expect improvements in this area as blade server technology develops.

Application Execution

Although application testing was not a focus for the blade PoC, we installed and ran a supply chain management application with which we were familiar. The database server for this application runs on an Intel® Itanium® processor, so for our test we ran the application components on the blade server and the database on the Itanium processor-based server.

We found no installation, configuration, performance, or compatibility issues while running the application in this configuration.

Summary of PoC Results

Our PoC tests identified no significant barriers to introducing blade servers into our environment.

Based on our PoC experiences, we came to the following conclusions:

- **Impact on data center.** We don't need to retrofit our data centers for an initial deployment of blade servers. We can follow typical data center strategies, such as hot and cold aisle placement of the blade servers. The data center raised floor will not need to be upgraded to support a full rack of six blade chassis. We will continue to monitor this situation as we deploy more blade servers into our environment and as blade server technology evolves.

Cable aggregation also will have a positive impact on the data center by decreasing the number of cables needed and the time to install the blade servers, reducing costs and simplifying deployment.

- **Key integration tasks.** Our PoC helped us identify several key areas of activity related to integrating blade servers into our existing infrastructure. These issues will be addressed in the engineering and production phases of deployment. For more details, see the Deployment Guidelines section later in this paper.

- **Reliability and stability.** A blade server system can provide significant redundancy at both the chassis level and the component module level. Although we have deferred extensive reliability and stability testing to the engineering and production phases of the project, we feel that this high level of redundancy will contribute to better reliability for blade servers compared to our existing rack-based systems.

Business Benefits

Table 2 on the next page shows some significant business benefits that blade servers provide, as supported by the PoC results. These benefits apply to our environment at Intel.

Depending on factors such as the data center architecture, the supported applications, and where a company is starting in the deployment process, business benefits may vary.

Table 2. Business benefits

Benefit	Estimated Improvement	Comments
Lower acquisition costs including SAN/LAN integration	10 to 20 percent in depreciation and network costs.	The cost benefits derive from the aggregation of ports and the sharing of components, such as management modules. For example, we anticipate a 40 percent savings in cabling costs for our particular infrastructure.
Lower service costs for deployment and troubleshooting and repair	2.5 hours to install and physically configure the blade system versus 8.5 hours for a comparable rack-mounted system.	We anticipate a decrease in service costs for troubleshooting and repair due to reduced complexity and aggregation of components.
Lower power, cooling, and space requirements	8.5 to 25.3 percent in data center facilities costs (15 to 25 percent savings in power and cooling).	Cost benefits derive from more efficient power utilization and increased density of computing resources.
Reliability improved	Possible “soft” savings in reduced site services and operations costs. Not measured during PoC testing.	We see opportunities for improvement in reliability due to the decrease in number of components, resulting in fewer places for failure and in chassis- and component-level redundancy reducing single points of failure.
Manageability improvements	Possible “soft” savings in operations and site services costs. Not measured during PoC testing.	A chassis-level shared view of assets simplifies resource management and makes it more efficient. (The manageability port cost is factored into the acquisition costs.)

Table 3. TCO analysis

Category	Estimated TCO Savings	Reason for Savings
Depreciation	10 to 20 percent	Reduced acquisition costs, including capital spending for servers, switches, cabling, and data center equipment and infrastructure
Operations	3.8 to 10.6 percent	Reduced costs for deployment and troubleshooting and repair
Facilities	8.5 to 25.3 percent	Lower power and cooling requirements in the data center as well as overhead expenses based on square footage and depreciation of the HVAC infrastructure

Table 3, on the previous page, lists the results of a TCO analysis for our environment. The following areas had no identified TCO benefit:

- **Software** – We primarily use enterprise licenses, so a move from rack-based systems to blade servers would have no effect on license costs.
- **Engineering** – We use a standardized build developed for our servers. Deploying blade servers would not make engineering of the standard build any easier or more difficult.
- **Maintenance** – Our existing maintenance model is based on maintenance contracts. Maintenance contract costs are based on acquisition costs and are accounted for in the acquisition depreciation.
- **Security** – We saw no changes to security that would affect TCO.

Deployment Guidelines

Our PoC helped us identify several key areas that will require activity to integrate blade servers into our existing infrastructure. These include:

- Developing a shared server acquisition and provisioning utility model. A shared utility model facilitates optimizing the use of blade server chassis components and minimizes acquisition costs through the sharing of network and SAN switches and management modules.
- Integrating the new Ethernet and Fibre Channel switches and new management capabilities into the existing infrastructure, which will require additional engineering testing and development of new operational procedures.
- Developing and implementing a new high availability model for failover of blade servers.
- For blade server environments with SAN, ensuring that the SAN boot is reliable and performs as expected since not all blade designs accommodate both on-blade redundant array of independent disks (RAID) and SAN at the same time.
- For blade server environments with SAN, dedicate a port on the blade chassis Ethernet switch for backup and recovery operations.
- Run the blade chassis Ethernet switch at 1,000 Mbps to mitigate potential network bottlenecks.

We are currently in the engineering stage of deploying blade servers into our environment. We will target rack servers that are reaching the end of their useful lifecycle for replacement by blade servers.

We will initially deploy blade servers for applications that require a maximum of two processors and have low storage requirements. Servers in this category include infrastructure servers, web servers, terminal servers, and low-end application servers.

Once we have engineered the functionality required to integrate blade servers into our existing SAN infrastructure, we will begin to deploy blades in environments that use SAN. Such an environment would typically use two-processor servers, up to 8 GB RAM per blade, and two NICs per blade, including file and print servers and low- to mid-end database servers.

Several additional factors to consider when integrating blade servers into an environment are:

- The life cycle of a blade server may differ from that of a rack server. Chassis transitions must be carefully managed using vendor roadmaps.
- Processor speed enhancements for blades may lag rack servers. As new processors come out, they will become available in rack-mounted servers first and later in blade servers. We estimate a 30- to 90-day lag in blade server processor speeds. In a business IT environment like ours, where a server life cycle may be as long as four years, a 30- to 90-day lag is not significant. However, it may be a more significant issue in a CPU-intensive environment, such as one supporting high performance computing (HPC) applications.
- Greater “lock-in” relative to rack servers. The lack of a standard blade chassis in the industry could lead to dependency on a single product.

Conclusion

Our PoC study explored the benefits and risks of integrating blade server technology into our existing infrastructure. Our Intel IT-wide team of experts explored a range of integration, installation, and power utilization issues. We found no significant barriers to introducing blade servers into our environment and a number of significant benefits to deploying blade servers.

Our TCO analysis showed a significant reduction of costs in the areas of acquisition and depreciation, operations, and facilities.

These benefits include lower acquisition and service costs, lower power, cooling, and space requirements, improved reliability, and improved resource management capabilities. Our TCO analysis showed a significant reduction of costs in the areas of acquisition and depreciation, operations, and facilities.

Based on our PoC study, we are preparing to deploy blade servers into our business IT computing environment.

Acronyms

CLI	command line interface
CPU	central processing unit
HBA	host bus adapter
HPC	high performance computing
KVM	keyboard, video, and mouse
NIC	network interface card
PoC	proof-of-concept
RAID	redundant array of independent disks
SAN	storage area network
SNMP	simple network management protocol
TCO	total cost of ownership
TTCP	test TCP
VLAN	virtual local area network
VPD	vital product data
WAN	wide area network

**For more information, visit our site on the World Wide Web:
www.intel.com/IT**

This document and related materials and information are provided "as is" with no warranties, express or implied, including but not limited to any implied warranty of merchantability, fitness for a particular purpose, non-infringement of intellectual property rights, or any warranty otherwise arising out of any proposal, specification, or sample. Intel assumes no responsibility for any errors contained in this document and has no liabilities or obligations for any damages arising from or in connection with the use of this document.

Intel, the Intel logo, and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.
Copyright © 2004 Intel Corporation. All rights reserved.

