

# CHAPTER 3

---

## The Virtualized Server Environment

Based on the analysis performed on the existing server environment in the previous chapter, this chapter covers the virtualized solution. The Capacity Planner run in the previous chapter produced a tremendous amount of data about the environment, taking the guesswork out of the drawbacks to the existing environment. It results in a more efficient new solution. Based on the statistics, the ideal virtualized server solution can be crafted with predicted results.

This chapter covers the new solution based on the Capacity Planner and some of the technologies integral to the solution, including blade technology. The last section covers blades in general, which you'll want to review if you're not familiar with blades technology.

### The New Server Environment

Capacity Planner produced a set of general server recommendations. These recommendations can be implemented in the form of standalone servers or blade servers. An overview of blade servers is presented later in this chapter.

Capacity Planner recommended servers with roughly the configuration shown in Table 3-1.

**Table 3-1** Recommended Consolidation Target Platform

Component	Number	Size, Speed, Type, Make and Model
CPUs (this means cores not processors)	8	3.2GHz Quad Core
Random Access Memory (RAM)		24GB
Network interface card (NIC)	6	HP NC364m Quad Port 1GBe
Host Bus Adapter (HBA)	2	Dual Channel Qlogic QMH2462 (Two embedded 1 GB NICs plus one quad port 1 GB NIC)
Local storage	2	72GB SAS RAID 1

As discussed in Chapter 2, several HP servers meet these requirements. To achieve maximum rack density and accommodate future growth, HP BladeSystem was selected as the target platform.

The ideal solution was 12 blades with the following parameters:

- 12 HP BL460c Blade Servers each with 2 quad-core processors (eight core total per blade)
- 24GB memory (upgradeable to 32GB)
- 6 NIC ports per blade
- 1 dual-channel Fibre Channel (FC) Host Bus Adapter (HBA) per blade

Figure 3-1 shows the BladeSystem proposed along with shared storage replacing the 134 individual servers.

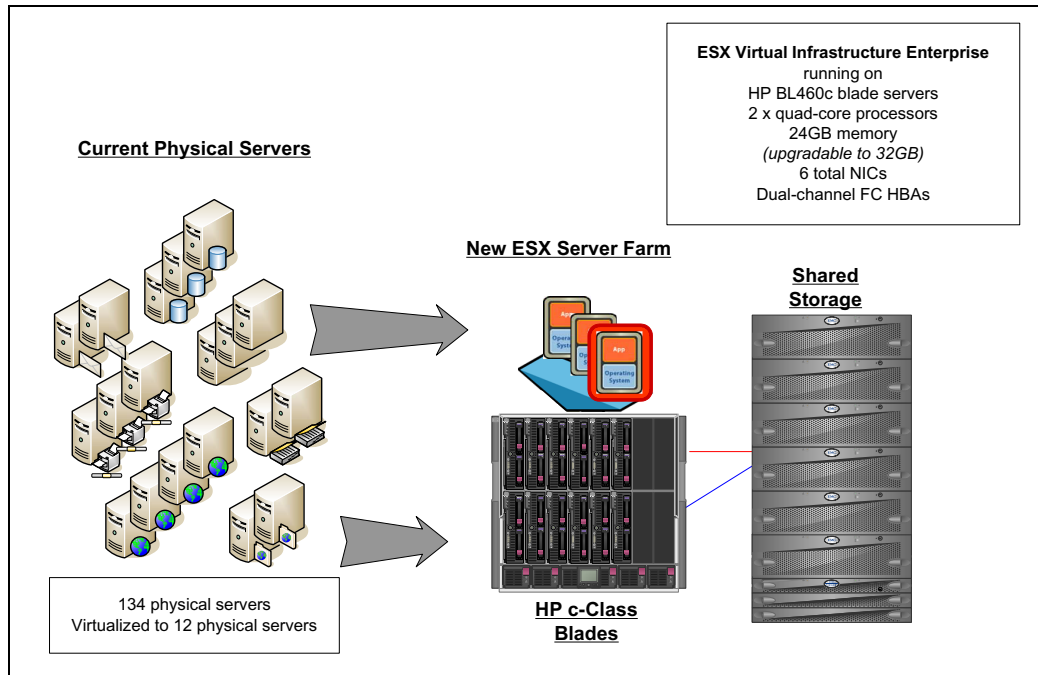


Figure 3-1 Virtualized Environment With Blade Solution

The savings the client will realize from reducing the 134 servers to 12 blades is immense, and the Return on Investment (ROI) is discussed in the next chapter.

The virtualized environment also requires shared storage, which partially existed before the virtualized design was crafted. The old shared storage was outdated, and the client determined that they needed new shared storage whether the virtualized environment was implemented or not.

To complete the savings estimate, I must calculate the power and cooling required to support the new environment. Table 3-2 covers the blade solution power and cooling.

**Table 3-2** Power and Cooling Breakdown of the Blade Solution

Model and Generation	Qty	Watts	Total Watts	Amps	Total Amps	BTU/hr	BTU Total
BL460c G5	12		5873		28.81		20029

The individual *Watts*, *Amps*, and *BTU/hr* are missing from Table 3-2 because the blade enclosure has power associated with it based on the specifics of the configuration, including the number and type of blades, the number of power supplies and fans, and other factors, such as I/O cards. The totals shown in Table 3-2 were painstakingly determined based on the specific client requirements and the BladeSystem configuration.

Figure 3-2 shows the HP Power Calculator for the BL460c G5 in a c7000 enclosure with full Ethernet and SAN connectivity.

Enclosure # 1, 10 Fans, 6 PSU				
Power Credentials	Idle	@ 100% Utilization	Circuit Sizing/HPTC	Unit(s)
Input Power	3302	5080	5873	Watt(s)
Input VA	3369	5184	5993	VA
BTU	11260	17323	20029	BTU(s)
Input Current	16.2	24.92	28.81	Amp(s)
Input Current Per Cord	2.7	4.154	4.802	Amp(s)
Airflow in CFM	256	389	583	
Airflow in CMM	7.239	11.018	16.517	

**Figure 3-2** HP Power Calculator for BL460c c7000 with Full I/O Capability

The power and cooling numbers for the blades compare favorably with the power and cooling from the 134 existing server solution, as shown in Table 3-3 and covered earlier.

**Table 3-3** Power and Cooling Breakdown of 134 Old Servers

<b>Model and Generation</b>	<b>Qty</b>	<b>Watts</b>	<b>Total Watts</b>	<b>Amps</b>	<b>Total Amps</b>	<b>BTU/hr</b>	<b>BTU Total</b>
DL360 G3	24	383	9192	1.9	45.6	1307	3168
DL360 G4	36	459	16524	2.3	82.2	1565	56340
DL360 G5	12	428	5136	2.2	26.4	1461	17532
DL360 G5	28	505	14140	2.5	70	1723	48244
DL380 G3	16	559	8944	2.7	43.2	1906	30496
DL380 G4	10	459	4580	2.3	23	1563	15630
DL380 G5	6	779	4674	3.8	22.8	2656	15936
DL380 G2	2	1043	2086	5.1	10.2	3557	7114
<b>Total</b>	<b>134</b>		<b>65276</b>		<b>324</b>		<b>222660</b>

I created power and cooling from both tables using the HP Power Calculator. The comparison between the existing and new virtualized solutions are compelling in terms of the power and cooling saved, shown in Table 3-4.

**Table 3-4** Comparison of Existing and Virtualized Environments

<b>Total 134 Server Power</b>	<b>Total 12 Blades Power</b>	<b>Blades Power Requirement Compared to Existing Environment</b>	<b>Total 134 Server BTU</b>	<b>Total 12 Blades BTU</b>	<b>Blades Cooling Requirement Compared to Existing Environment</b>
65276	5873	9% of existing environment	222660	20029	9% of existing environment

A savings of 90% in both power and cooling for servers is compelling from an efficiency standpoint. Many other factor make up payback, such as floor space, but from purely a server standpoint, you can't do much better than that.

Because of the immense savings in power and cooling, this client decided to double the blades environment to accommodate the planned growth of roughly 100 servers over the next year. Figure 3-3 shows 12 blades in a rack, which replaces the existing server environment and will be doubled to support the additional 100 servers over the next year.




Figure 3-3 Virtualized Environment

Figure 3-3 shows how little space in the rack is consumed by the 24 blades. There is a substantial amount of space left for additional components. This huge savings in floor space is part of the ROI covered in Chapter 4 along with the many other savings realized by the blades-based virtualization solution.

Capacity Planner uses numerous assumptions, which are covered in the next section.

## Assumptions

Some key assumptions are produced in *Capacity Planner* in order for the new, virtualized environment to operate:

- All target VMware ESX host servers will be identical in configuration. Differences in configurations can impact sizing and may require modeling additional scenarios.
- All planned hosts will be located in two physical locations.
- All NIC ports in the target ESX hosts will be active. If not all ports are active, more ESX hosts may be required to distribute the potential network load to avoid network-based performance bottlenecks.
- The use of Gigabit NICs assumes a network using Gigabit speeds.
- Shared storage using a SAN will be available, which, in this specific client example, required an upgrade to the SAN.
- All target ESX hosts will have access to required storage, and storage limitations are not considered a gating factor.
- An additional server and database will be required for the Virtual Center management server.
- All workloads will utilize VMware virtual disks residing on a Virtual Machine File System (VMFS), and no virtual disks will map directly to SAN Logical Units (LUNs).
- Servers considered for reuse will have processor specifications to support VMotion compatibility. Non-identical hardware will limit the ability to use VMotion and restrict the ability to shift loads across multiple ESX hosts. VMotion allows you to move a virtual machine from one physical server to another with no impact to users.



- The target ESX hosts are reserved exclusively for server consolidation of the existing servers and server containment to support future provisioning requests.
- All target ESX hosts will be utilized to their maximum potential and workloads can be freely balanced across the target hosts with no network limitations.

These assumptions are key to designing the target solution. Select assumptions can be changed; however, this may result in a modified solution and affect the ROI.

The next section covers some rudimentary background in blades.

## Blades Background

A blade is a server in a format that allows it to be inserted into an enclosure housing numerous blades. The enclosure provides some common services, such as power, cooling, management, network, and SAN interconnects. The blade has its own resources, including processor(s), memory, disks (optionally,) NICs, and other components. Each blade requires its own operating system, just like a rack-mounted or tower server, and is used in a way similar to any other physical server.

The blades used in the virtualized solution presented in this chapter are part of the HP c-Class BladeSystem. Some of the c-Class advantages include the following:

- Smaller physical format saves rack space in the data center. In the case of the HP c-Class c7000 enclosure, up to 16 half-height blades can be housed in 10U of rack space, including all power and interconnecting switches. A special-purpose blade that provides 32 independent server nodes in 10U of space is also available. Ethernet and SAN switch modules can be optionally added to the enclosure saving additional space compared to top-of-rack switch alternatives.
- Power and cooling are greatly reduced with blades. The results are higher power efficiency, dynamic power scaling, and a subsequent reduction in heat output. Air circulation is also optimized to efficiently exhaust the heat that is produced.

- HP's management suite allows management at the enclosure level and at the server blades. Shared and individual resources at the blades are all controlled and reported on through a common interface. Software deployment, patch management, and power management are some of the tools supported by HP's Integrated Control Environment (ICE) that's part of the BladeSystem. No separate connections to keyboard, video, and mouse are required, because Integrated Lights Out (iLO) allows full graphical remote control from any WAN-attached device with a browser.
- With common power, network, SAN, and management interfaces, cabling can be tremendously reduced. A fully provisioned rack-mount server could require 2 power cords, a keyboard, monitor, and mouse, up to 6 network connections, an iLO management connection, and 2 SAN connections. Those 14 cables mean 224 total connections for 16 conventional servers. In a c-Class blade environment, these needs could typically be met with 6 power cables, 2 iLO management connections, as few as 6 total network connections, and as few as 2 total SAN connections, resulting in a 14-to-1 reduction in cabling.

Figure 3-4 shows a front view of a c7000.



Figure 3-4 HP BladeSystem Front View

Figure 3-4 shows eight half-height and four full-height server blades and power supply modules, which are hot-pluggable from the front. An integrated touch-screen display allows easy access to status information and speeds up the initial configuration process.

Figure 3-5 shows a rear view of the c7000.



Figure 3-5 HP BladeSystem Rear View

Hot-pluggable fans plug in from the rear of the enclosure as well as single phase or 3-phase power. Management and connection to the iLO network interface is provided through one standard and an optional redundant Onboard Administrator Module.

Eight interconnect bays give flexibility to provide the required level of Ethernet, SAN, InfiniBand, and SAS storage connectivity. These can be simple pass-through modules (similar to patch panels) or full-function switch modules that reduce or eliminate the need for infrastructure or top-of-rack switches.

HP Virtual Connect modules give an option for Ethernet and SAN connectivity that allows all resources to be defined once at the start of the process, with dynamic provisioning of blades through the life of the enclosure.

All the interconnect modules are typically installed in pairs, and map to the physical NIC ports and mezzanine cards installed on each Blade server. Everything is tied together through a signal midplane and a power backplane.

The virtualized solution in this chapter consists of only server blades; however, storage blades, tape blades, and PCI expansion blades are also available.

Server blades are available in configurations from a single processor to 4 quad-core Intel or AMD processors in half-height and full-height formats. A special-purpose blade combines 2 independent server nodes in a single half-height blade carrier, allowing up to 32 servers in a single c7000 enclosure. Itanium processors are also supported in two current full-height and double-wide full-height Integrity Server Blade models.

Storage blades are available to provide either local direct-attached drives for an adjacent server blade or common access to other blades through iSCSI or NAS connectivity.

Tape blades give the ability to run tape backups within the same physical blade enclosure.

PCI expansion blades allow industry standard PCI-Express or PCI-X cards to be utilized by an adjacent server blade.

More information is available on c-Class blades at [http://h18004.www1.hp.com/products/blades/components/c-class-components.html?jumpid=reg\\_R1002\\_USEN](http://h18004.www1.hp.com/products/blades/components/c-class-components.html?jumpid=reg_R1002_USEN) at the time of this writing.

## Performance Analysis After Virtualization

After virtualization has been implemented, there are specialized tools to assist with performance analysis. After virtualization has been implemented, analysis is more complex because there are many virtual machines running on a single server. One such tool my firm has used on virtualized environments is Akorri BalancePoint. This performance-analysis tool for servers, storage systems, and applications is “virtualization aware” and ideally suited for use with VMware. Because of the virtualization of all the normally physical components of a data center—network, server and storage—it can be difficult to diagnose and resolve problems with traditional software tools. BalancePoint does the following in your virtualized environment:

- Displays a mapping of physical to virtual resources for each application running on a VM
- Measures the performance of physical and virtual elements
- Provides alerts and recommendations when problems are discovered
- Provides historical baselines of all virtual machines
- Identifies which virtual machines are overutilizing resources
- Provides detailed I/O data
- Provides comprehensive reporting

Even the most well-run data center is not immune to problems and bottlenecks. BalancePoint can provide invaluable troubleshooting capabilities. It facilitates deep analysis into the critical elements: operating systems, databases, servers, storage, virtualization systems, helping to accelerate the troubleshooting process. Figure 3-6 is a performance overview produced by BalancePoint.

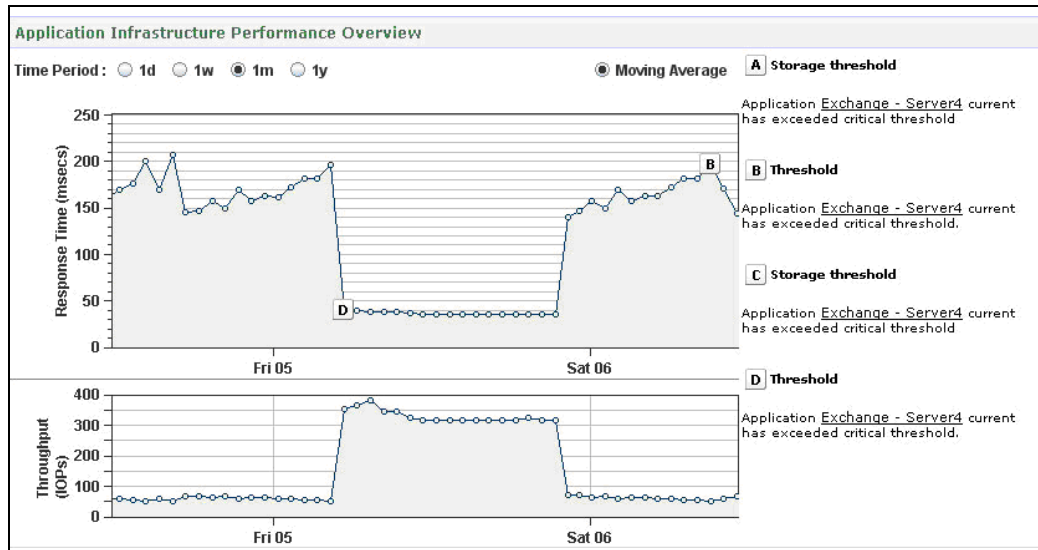


Figure 3-6 Performance Overview

The report shown in Figure 3-6 illustrates the response time and throughput of a Microsoft Exchange server with the underlying storage system. The thresholds can help alert potential bottlenecks. Response time is often the measure of whether an environment has been successfully virtualized, so this detailed analysis is immensely helpful.

Despite the gains in server utilization offered by virtualization, capacity planning is still a challenge. BalancePoint helps indicate where additional resources are available and where more virtual machines can be allocated without running into capacity or performance problems. Rather than waiting until problems are evident, BalancePoint can help an administrator strategically plan how virtual and physical resources are allocated. Virtualization is addictive, and with these useful outputs, you can clearly see where you can add additional virtual machines and still maintain the desired response times. Figure 3-7 shows the response time.

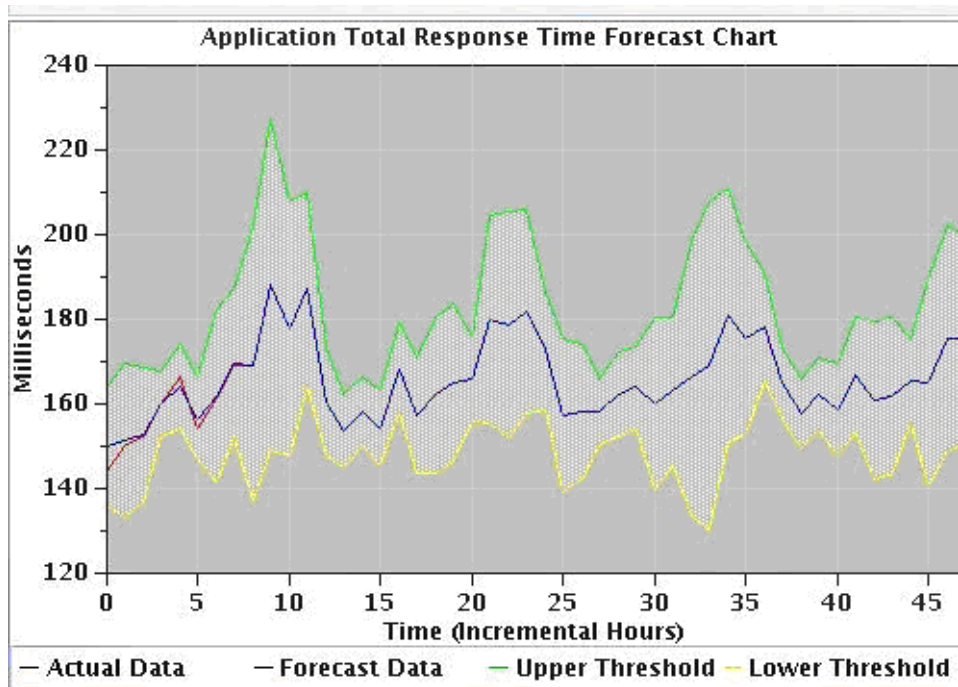


Figure 3-7 Response Time Forecast

Although the colors in Figure 3-7 are lost, you can see that actual and forecast data are provided. The tool shows what you can expect to see when applications are added to the environment which, in some cases, may have an adverse impact on the applications currently running on the server.

It is easy to see how this tool can be invaluable for managing a virtual infrastructure. This, and similar software tools, can make the already impressive ROI that virtualization provides even more impressive with the ability to add forecast data to existing data.

## Summary

The benefits presented in this chapter make a compelling argument for server virtualization. The results presented are typical for server virtualization. The case study illustrates what can be expected in any server



virtualization and is by no means extreme. Every such assessment I've run has resulted in a server virtualization project because, with immense benefits presented in this chapter, it is difficult to decline taking on server virtualization.

The next chapter covers the Return on Investment (ROI) for the new environment.

