FROM VISION TO REALITY — ONLINE LEARNING IN A COMPLETELY DIGITAL WORLD

A SCALABLE REFERENCE ARCHITECTURE FOR BLACKBOARD LEARN™

January 2010

Contents

Tables

Figures

Chapter 1

Introduction

A paradigm shift is underway in online learning. Online learning tools that were once used primarily to augment classroom learning have matured to support a fully online model in which the Internet serves as the primary medium for teaching and learning. As the prevalence of high-speed Internet connections and home computer ownership has grown, so too has the demand for distance learning. Today's learners are turning to Web-based distance learning programs as a means to balance other obligations with the desire to further their education, or because they are too far away from the institutions that can best serve their needs.

Research conducted by the Blackboard Institute (www.blackboardinstitute.com/online_programs.asp) indicates that today's educators also see advantages to delivering courses online. Online learning programs can enable institutions to:

- Provide greater access to teaching and learning while expanding into new audiences of students
- Meet the needs of non-traditional students
- Overcome physical space limitations at an academic institution while supporting institutional growth
- Develop alternative revenue streams

According to the 2008 Sloan Survey of Online Learning, online enrollments have continued to grow at rates far in excess of the total student population for higher education. The trend toward online learning has given rise to a number of U.S. institutions where courses are offered almost exclusively online, including Capella University, Grand Canyon University, Walden University, University of Phoenix, Kaplan University, and National University. Many of these institutions have grown rapidly in the last decade. For example, National University had 480 students enrolled in its online classes in 1999 and now has almost 14,000 students participating online. The University of Phoenix, the largest private university in the U.S., now has more than 63,500 online students, and realized an 80 percent growth between 2001 and 2002.

The movement toward these fully online learning models creates new demands for the IT environment and opens up new opportunities for institutions to use technology as a competitive tool. Demand for higher performance and faster response times are being driven by two key factors: greater numbers of users, and increased activity by users. Not only are institutions growing and supporting greater numbers of online users, but users are also interacting with the system in new ways that generate substantially greater system load. Users tend to stay online for much longer periods of time and have greater interactivity with their online courses. The technology must therefore support large numbers of users who concurrently access the system and require rapid response times so that they can interact without noticeable delays.

When users have consistent and responsive page load times, online learning is easier and more effective, New forms of teaching and learning can also evolve when interaction has the feeling of being real-time. Traditional online lectures and online documents can be blended with interactive media, allowing students to share experiences with each other across geographical boundaries. This can increase the effectiveness of the learning environment and expand the available market the institution serves. Institutions want their online learning technology infrastructure to support these new forms of teaching and learning on a scale that allows many more students than could traditionally participate in a classroom approach.

The online learning architecture must therefore be highly scalable as well as cost-effective to maintain so that institutions can focus on teaching instead of technology. The right technology infrastructure for online learning can create a competitive advantage for the institution.

Chapter 2

Dell Reference Architecture for Blackboard Learn™

Dell, Blackboard, Citrix, and Quest have come together to define a scalable reference architecture that takes advantage of today's latest technologies to support a rich online learning or distance learning model while offering superior performance and scalability as well as predictable service levels. The Dell reference architecture for Blackboard Learn™ is a blueprint for successfully optimizing the deployment and management of Blackboard software on Dell servers and storage. It identifies key system components and best practices that are required to achieve high service levels and scalability along with efficient IT operations.

The reference architecture includes the following:

- An optimized configuration for Blackboard Learn software with Citrix® XenServer™ software and Red Hat® Enterprise Linux® 5.
- Robust hardware server and storage architecture based on Dell EqualLogic™ iSCSI storage systems and Dell™ PowerEdge™ servers based on the latest Intel® Xeon® 5500 series (Nehalem) processors.
- Optimization and tuning of the Oracle 11g database with Blackboard Learn.
- Best practices for deploying the next generation of Blackboard Learn.
- End-to-end monitoring and management using management tools from Quest Software to provide the visibility needed to optimize the overall environment and manage it efficiently.

The reference architecture underwent rigorous testing with a team of engineers from Dell, Blackboard, Citrix, and Quest who also designed the architecture. The purpose of the testing was to validate the architecture and fine-tune its design for maximum scalability. Testing included performance benchmarks, which affirmed that the reference architecture could support in excess of 100,000 active sessions when deployed with an application tier that consisted of three Dell™ PowerEdge™ R710 servers running Citrix XenServer software as the virtualization platform along with a fully loaded Dell PowerEdge R900 server as the database server. Figure 1 shows a logical representation of the reference architecture with different tiers for the application, database, storage, and management components of the architecture.

The architecture is designed to support maximum scalability and high service levels through virtualization and an optimized management environment. It enables new levels of scalability while providing a very cost-effective modular solution. This enables institutions to start with a cost-effective small pilot and then grow and expand over time to support an enterprise-scale online learning environment.

The reference architecture also includes best practices for deployment and management of each tier as described in the remaining chapters of this document. A complete list and descriptions of the recommended components in the Dell Reference Architecture for Blackboard Learn is included in Appendix A.

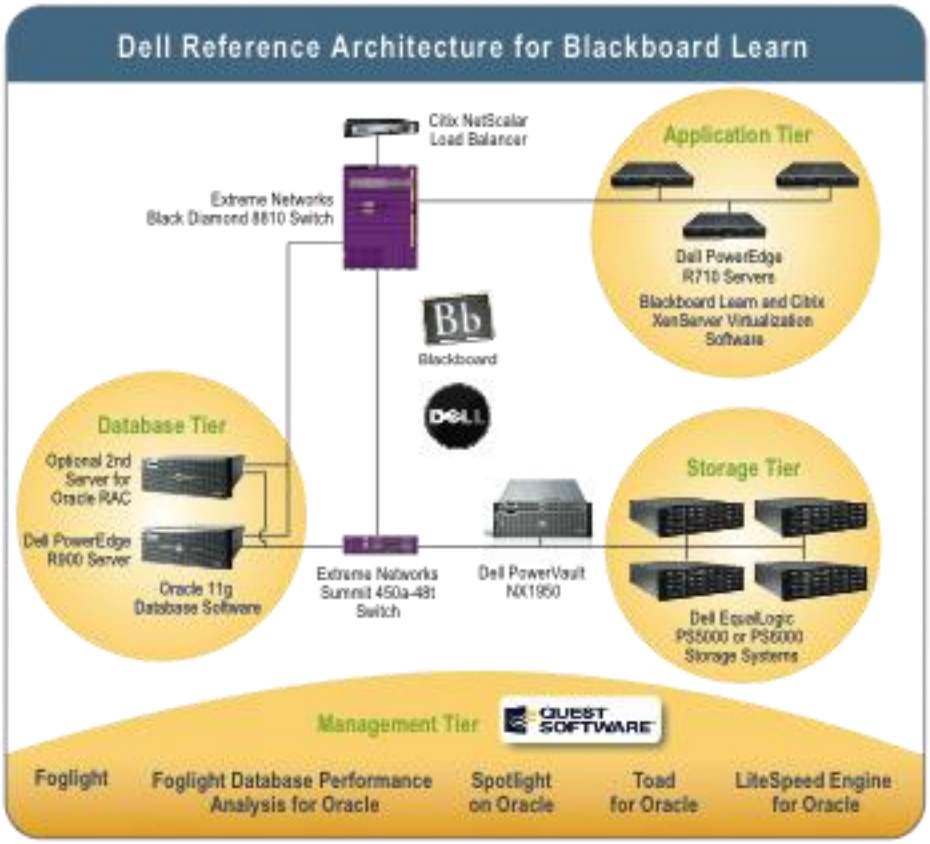

Figure 1: Logical Diagram of Dell Reference Architecture for Blackboard Learn

Flexible and Scalable Application Tier

Application tier scalability is achieved using a highly flexible horizontally scaled infrastructure. Multiple instances of the Blackboard Learn application can be deployed on a single Dell server using the Citrix XenServer virtualization platform to isolate Blackboard application instances from each other and to manage allocation of CPU and memory resources.

Each Dell server in the application tier runs multiple virtual machines (VMs), each with its own instance of the Blackboard Learn software. To achieve a high volume of concurrent user sessions, the architecture employs load balancing to distribute requests across all of the Blackboard Learn application instances deployed on the virtualized servers.

The architecture also takes advantage of 64-bit computing by using a 4GB Java heap for each Blackboard application instance. The Java heap had been limited to 1.7GB in a 32-bit environment, but this limitation is overcome in the 64-bit environment with Red Hat Enterprise Linux and Dell servers based on the 64-bit Intel Nehalem processor. The architecture includes a recommended 8GB of memory for each VM, leaving 2GB for the Apache Web server and 2GB for the OS and monitoring tools. The larger memory footprint for Blackboard Learn enables each application instance to more efficiently service a high volume of user requests. In other words, each application instance can scale vertically while the architecture can also be scaled horizontally by deploying additional VMs. This best of both worlds approach takes advantage of many application threads in the application tier to service thousands of requests per minute.

The recommended servers for the application tier are Dell PowerEdge R710 servers, which were chosen for their small footprint, low cost, and powerful performance with Intel 5500 series processors. The application tier could also be built using other Dell servers such as Dell™ PowerEdge™ R610 rack mountable servers or Dell™ PowerEdge™ M610 or M710 blade servers, all of which feature the latest Intel 5500 series processors.

The highly flexible and scalable application tier can help customers achieve:

- Lower costs – Increased utilization of servers enables good performance on a low-cost consolidated infrastructure.
- Increased flexibility – Virtual machines can be easily moved to other physical servers to redistribute workloads or recover quickly from a hardware failure.
- Faster Deployment – With fewer physical systems to setup and configure, less time is required for deployment.

Virtualization for Higher Utilization

The design of the application tier is based on Citrix XenServer technology, which is optimized for Dell PowerEdge servers and Dell EqualLogic storage solutions. Dell and Citrix have partnered to deliver pre-qualified virtualization-ready platforms for today's dynamic and growing data center needs.

Citrix® XenServer™ Dell Edition is a powerful virtualization solution based on the 64-bit open-source Xen® hypervisor. It enables multiple instances of the Red Hat Enterprise Linux operating system to be deployed on each physical server. Server processor, memory, storage and networking resources are shared across multiple virtual machines, yet each virtual machine and operating system instance can have direct control over specific system resources. Blackboard Learn software can then be installed and operated on each VM in much the same way as running each instance on a separate physical server.

There are two key reasons for implementing virtualization in the application tier. The first is that modern CPUs such as the Intel 5500 series processors are so powerful that a single instance of the operating system and its application software components typically cannot fully utilize the CPU. It is more cost-effective to add additional memory to a server and divide the server resources into multiple virtual machines than it is to buy more servers. Secondly, virtualization greatly simplifies provisioning and management of the Blackboard Learn environment. Citrix XenServer software includes tools to enable administrators to quickly and easily clone an image of a virtual machine, including the operating system and all application instances. This cloned VM image can then be quickly provisioned to a new VM when, for example, greater capacity is needed. Cloning also helps ensure that all instances of the application are deployed in identical configurations, greatly simplifying ongoing maintenance of the environment.

Citrix XenServer operates from within the internal flash storage in Dell PowerEdge servers, thus enabling high performance and efficient utilization of system resources. Citrix XenServer also includes advanced monitoring and management capabilities that help simplify management of the virtualized environment.

Applications can be cloned by making copies of an existing VM or application server. The results can then be stored as templates or exported as VMs. Using a template or exported VM is a simple way to make a master copy from which additional copies may be provisioned. The steps for cloning or exporting this master copy of the VM can be found in the XenServer Guest VM Guide at support.citrix.com/article/ctx121750.

Citrix XenServer is a free virtualization platform based on the open-source Xen® hypervisor and includes XenCenter™, a multi-server management console with core management features such as multi-server management, virtual machine (VM) templates, snapshots, shared storage support, resource pools and XenMotion™ live migration. In addition, Citrix offers advanced management capabilities in its Citrix Essentials™ for XenServer™ product line. Citrix Essentials for XenServer is available in two editions, Enterprise and Platinum.

XenServer is truly a bare metal solution. Built on the open-source Xen hypervisor, it uses a combination of paravirtualization and hardware assisted virtualization. Paravirtualization allows a guest OS to be fully aware that it is being run on virtualized hardware, allowing for the simple, lean hypervisor to deliver much higher performance. XenServer is designed to leverage hardware virtualization assist technologies such as Intel® Virtualization Technology (Intel® VT). It also leverages the native storage capabilities of Dell EqualLogic storage systems as well as other native storage facilities such as NFS, Microsoft VHD, iSCSI, and Fibre Channel and SAN technologies.

Highly Scalable Web Delivery

User requests in the virtualized environment of the Dell reference architecture for Blackboard software are load-balanced with the Citrix® NetScaler® Web application delivery controller. The controller functions as an application accelerator with caching and HTTP compression. It provides advanced traffic management through load balancing and content switching functions for Layers 4-7 of the network stack.

Citrix NetScaler can multiplex a large number of client side connections into a small number of back-end server connections that are “long-lived” sessions, offloading most of the laborious session set-up/tear-downs associated with TCP/IP. On the client side, more than five million simultaneous connections can be maintained, enabling extensive Blackboard deployments with hundreds of classes and thousands of students running multiple sessions each.

NetScaler is available either as a purpose-built physical appliance or as a virtual appliance that can be hosted as a XenServer VM, enabling it to execute on the same Dell PowerEdge servers as the Blackboard Learn software.

Further information about the benefits of using Citrix NetScaler with Blackboard Learn can be found at citrix.nl/site/resources/dynamic/accessAnswers/ENHANCEBLACKBOARD.pdf. Recommended configuration settings for Citrix NetScaler with the Blackboard Learn environment are described in Chapter 4, Best Practices and Configuration Guidelines.

Storage Tier with Pooled Resources

The reference architecture provides a very flexible storage environment that takes advantage of integration between Citrix XenServer software and Dell EqualLogic PS Series storage systems. Dell EqualLogic PS Series storage systems are recommended because they offer affordable enterprise-class iSCSI storage that provides both the high performance of iSCSI and the ability to increase availability in a SAN infrastructure.

EqualLogic PS Series arrays simplify storage deployment by offering seamless access to a single pool of storage that can be easily allocated to specific components of the reference architecture. The arrays also include intelligent automation of storage management. They offer single console management and ease of storage provisioning to increase the power and flexibility of the storage infrastructure. Additional arrays can be seamlessly added into an existing SAN to automatically increase storage pool resources without disruption of the Blackboard Learn application or the availability of its data.

Benefits of the Storage Tier Design

The storage infrastructure based on Dell EqualLogic PS Series storage systems and Citrix XenServer software brings the following unique benefits to the Dell Reference Architecture for Blackboard Learn:

- **Optimized utilization** – Storage utilization is optimized by balancing loads dynamically across multiple storage arrays as usage needs change.
- **Reduced deployment time and effort** – The EqualLogic PS Series Group Manager enables administrators to easily create new volumes and assign them to XenServer hosts. The XenServer EqualLogic Storage Adapter also enables storage volumes to be created from the XenCenter interface or using XenServer CLI.
- **Reduced management complexity** – EqualLogic PS Series SANs simplify storage management by consolidating physical storage and providing a single-pane management view of the entire virtualized storage pool. Storage volumes for virtual machines can also be managed directly from XenCenter or the XenServer CLI.
- **High availability** – In addition to the best-in-class reliability features offered by Dell EqualLogic PS Series arrays, Citrix XenServer 5.x introduces an automated local high-availability feature that protects virtual machines against host failures by restarting failed hosts on remaining servers in the XenServer pool. NIC bonding available in XenServer also protects against failures in either network card.
- **High scalability** – Adding more PS Series array members to a PS Series group allows storage capacity to scale along with performance and does not disrupt application or data availability. Using XenServer Resource Pools, administrators can also add new XenServer hosts to an existing pool to help meet increased workload demands.
- **Increased flexibility** – The net result of all these features and capabilities is to enable IT administrators to respond quickly and flexibly to enterprise demands and initiatives.

Flexible Storage Architecture

As shown in Figure 2, the PS series storage systems are presented as a virtualized storage pool that serves the entire Blackboard application while dedicating specific storage resources to specific components of the reference architecture. Separate disk volumes can be configured to provide storage for the database, the Blackboard content file system, application binaries, and temporary storage for virtual machines.

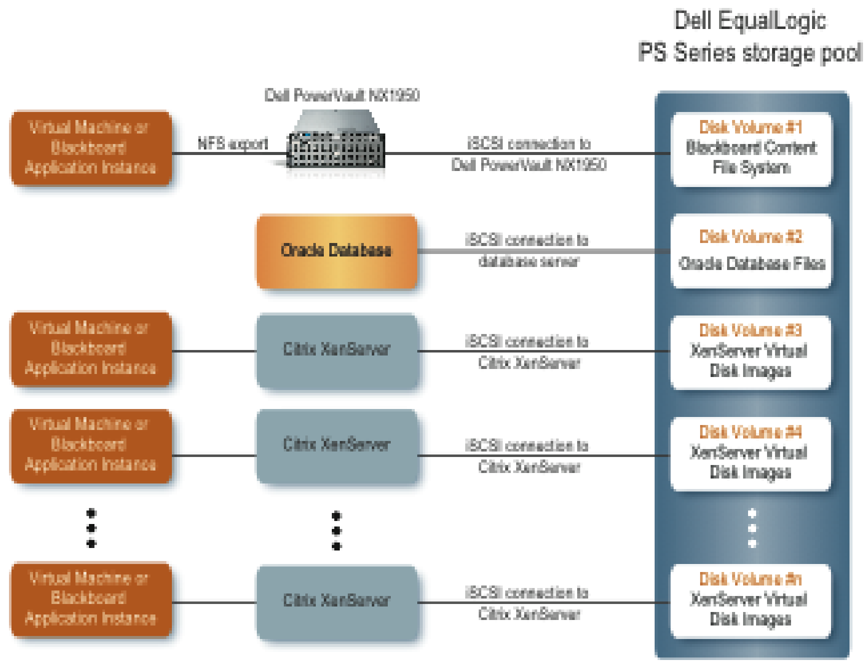

Figure 2: Flexible storage architecture based on Citrix XenServer

Different disk volumes are used for different types of storage so that each disk volume has a consistent type of storage that can be used for a specific application. For example, the PS5000E or PS6000E storage system can be populated with SATA drives that provide cost-effective storage for a file system. These systems are used in this reference architecture to store the content file system within Blackboard Learn. The PS5000E/PS6000E storage system and the Blackboard content file system are represented by disk volume number one in the figure.

An NFS export through the Dell™ PowerVault™ NX1950 system enables the Blackboard application to access this content file system via an NFS mount. The PowerVault NX1950 system communicates to the PS5000E storage system via the iSCSI protocol and then presents disk volume number one as an NFS file system available the application tier.

All of the other disk volumes in the reference architecture reside physically on PS5000XV storage systems, which contain high-speed (15K RPM) SAS drives. Disk volume number two represents storage for the Oracle database. The SAS drives in this disk volume provide the high performance needed for database reads and writes.

The remaining disk volumes in the figure (volumes three through n) are managed by the Citrix XenServer hypervisor and provide temporary storage for the application tier as well as a central location for storing application binaries used during startup. These disk volumes consist of virtual disk images that are defined in Citrix XenServer and then assigned as temporary storage for specific application instances running in a virtual machine. The additional level of abstraction provided by the virtual disk images enables greater flexibility for dynamic load balancing and allows the Blackboard Learn application instances to be easily moved between XenServer hosts.

Each Citrix XenServer host can use multiple disk volumes simultaneously, including different types of physical storage. This virtual storage approach also allows the VM host to take advantage of advanced storage features such as sparse provisioning, VDI snapshots, and fast cloning on Dell EqualLogic PS Series storage systems.

The only item in the reference architecture that is not stored in this centralized EqualLogic PS Series storage pool is the Citrix XenServer hypervisor. The XenServer hypervisor is installed locally on each Dell PowerEdge R710 server.

Further information about Citrix XenServer and its integration with Dell EqualLogic storage offerings can be found in the white paper titled, "Deploying Citrix XenServer 5.0 with Dell EqualLogic PS Series storage." This paper is available for download at www.citrix.com/site/resources/dynamic/partnerDocs/Citrix-XenServer-Equallogic_final.pdf.

Database Tier

The database tier is a vertically scaled architecture that can also include multiple database servers configured in an Oracle RAC environment. The Oracle database requires significant vertical scalability in the server hardware platform because a single large database instance is used by all of the Blackboard Learn application tier components. Using Oracle RAC enables a cluster of multiple physical servers to appear as a single large database server to the application tier. Oracle RAC should be considered in cases where extreme database scalability is required, or where application availability needs exceed that which can be expected from a single database server without a failover node.

The database server must be sized based on the number of database connections expected from the application tier. The more VMs contained in the application tier, the greater number of simultaneous database connections are possible. Each database connection requires CPU and memory resources, so the server must scale to support the expected workload. Chapter 4 provides a detailed explanation about sizing the database server.

The Dell PowerEdge R900 server, Dell's first 24-core Intel server, provides an excellent platform for the vertical scalability needs of the database server, enabling Blackboard Learn to be scaled to new levels with a single database server. This powerful server offers the following advantages:

- Proven application performance with Blackboard Learn.
- Excellent price/performance.
- Ability to optimize the need for expansion capacity versus rack density.
- Plenty of capacity for future growth with up to 24 CPU cores.
- Enterprise-class reliability features to help protect against unwanted downtime.

Depending on the performance requirements and the complexity of the customer database, some tuning of the Oracle database may be required. Chapter 4 provides an overview of best practices including recommendations for optimizing Oracle database performance.

Network Environment

Due to the high volume of data that flows between servers and storage systems in the reference architecture, a high-performance Ethernet network is recommended. The storage tier should also be configured for a storage area network (SAN) to enable high throughput and reliability.

The network architecture is based on network components from Extreme Networks. Extreme Networks provides converged Ethernet networks that support data, voice, and video for enterprises and service providers. The company's network solutions feature high-performance and high-availability switching that delivers insight and control, enabling customers to solve their real-world business communications challenges. Extreme Networks solutions are complemented by global, 24x7 service and support.

The recommended switches from Extreme Networks (BlackDiamond 8810 and Summit X450a-48t) are described in Appendix A.

Chapter 3

Validating and Optimizing Performance

Performance tests were conducted to validate the components of the reference architecture in order to optimize for high-performance. Multiple configurations were tested to determine the optimal combination that would support the greatest scalability. This chapter describes the benchmark process and interprets the performance results in terms of supporting a Blackboard configuration of this magnitude. Several best practices regarding setup and configuration of the environment were obtained through the testing process and are also documented in the final section of this chapter.

The Engineering Team

The benchmark was performed by a core team of engineers from Dell, Blackboard, Citrix, and Quest Software as well as a database specialist from Method R Corporation. The tests were performed at a Dell testing facility in Round Rock, Texas using servers and storage systems provided by Dell. Citrix and Quest Software also provided their respective hardware and software components for the benchmark tests. Quest Software tools were used throughout the tests to monitor the environment and provide data that would help identify the source of performance bottlenecks.

The engineer from Method R provided tools for profiling activity in the Oracle database and worked with the team to identify the most important business operations to optimize. Method R Corporation is a software tools company that offers consulting services for optimizing software applications built on the Oracle Database. Method R has partnered with Blackboard on multiple performance, scalability, and architecture projects and their methodology and profiling tools are commonly used by Blackboard. Software tools such as the Method R Profiler and the MR Tools suite can help DBAs, performance analysts, and application developers solve and prevent application response time problems.

Benchmark Process

In order to develop testing criteria that accurately reflect the needs of large online learning institutions, Blackboard sampled customer data sets that met this profile. A data model was created to match customer criteria such as the volume of data that is typically stored, how long it is retained, the rate of data creation, and the structure of data. After studying the data patterns, it was determined that a data model containing 25,000 course sections (5,000 active) and 800,000 users (125,000 active users) would be a good representation of the profile.

Blackboard also sampled user activity patterns to determine typical user loads and the typical activities of remote users during peak periods such as near the end of a school term. Based on the findings, a sample test case was created that mimics a typical user load. The breakdown of actual activities in the test case is identified in Appendix B.

It is worth noting that some functional areas such as exams and quizzes require more system resources than others. Therefore, it is also important for customers to identify the breakdown of how their users tend to utilize Blackboard Learn. If the breakdown is substantially different from the usage patterns described in Appendix B, results will vary. Some adjustments may therefore be needed when sizing the hardware and software configuration based on the performance results and sizing guidelines presented in this chapter.

Benchmark Results

One of the key goals of the performance optimization effort was to maximize the number of active sessions within a predefined response time threshold. An active session is defined as an authenticated user who is currently logged into the system. While each user may not be actively using Blackboard Learn during their entire login time, a session is still considered active as long as the user is logged in.

The maximum supported load for a given configuration was determined by monitoring several thresholds including user response times and system resource utilization metrics. User response time was considered the most important metric and the workload was considered to be within the bounds of the desired service level when all of the following conditions were met:

- 99% of user responses within five seconds
- 95% of user responses within three seconds
- 90% of user responses within one second

If the user response time exceeded any of these thresholds, the test run was rejected as not qualifying. Tuning and optimization work was performed to increase the number of transactions that could be achieved while staying within these bounds on user response times.

Benchmark tests revealed that optimum performance could be achieved when each PowerEdge R710 server was virtualized into four VMs with 8 GB of memory allocated to each VM. The PowerEdge R710 servers were configured with two quad-core Intel® Xeon® X5570 2.93 GHz processors with hyperthreading, which has the effect of making the eight cores resemble 16 cores. Hyperthreading in the Intel Xeon X5500 series CPUs typically yield good results for applications such as Blackboard Learn where there are many small transactions that are all independent of one another so they can be processed in any order. (Note that results may vary with a different CPU model or different clock speed.)

Figure 3 shows that the system showed near linear scalability from four to twelve VMs in the application tier. Each XenServer VM represents a single instance of Blackboard containing of an Apache web server and one 4GB Tomcat application server. The Quest Foglight monitoring tools were embedded within the XenServer VM and the Blackboard instances.

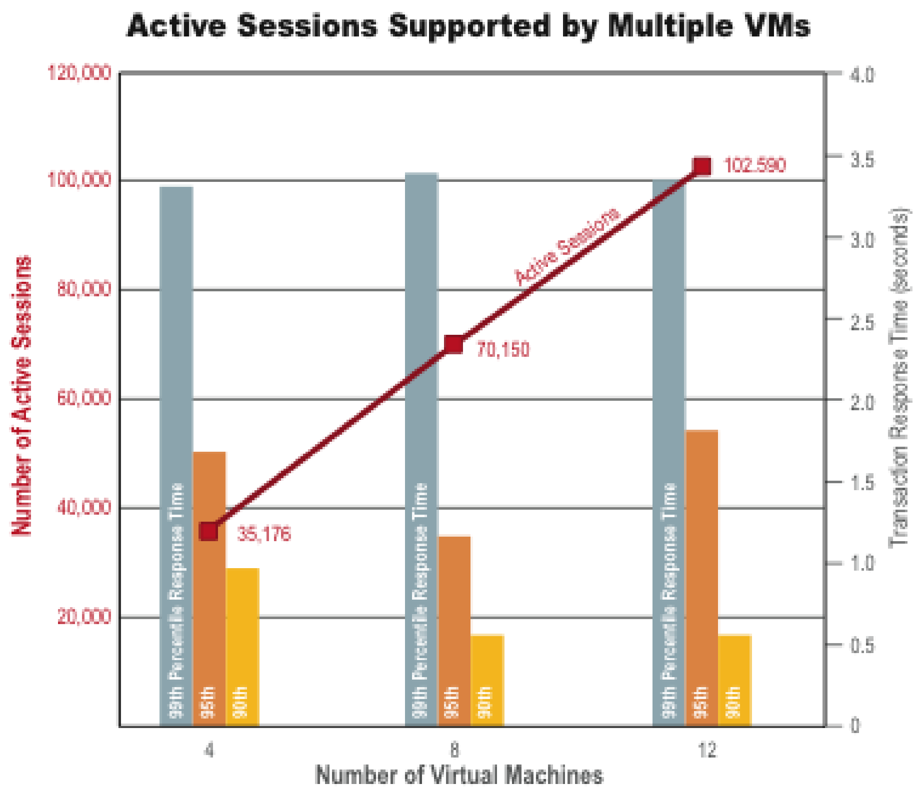

Figure 3: Scalability to more than 102,000 active sessions with twelve VMs

The total number of active sessions reached more than 102,000, and the system supported consistent throughput of 2900 hits/second in the 12 VM configuration. This total scalability exceeded the expectations of the benchmark team and shows that a large user population for Blackboard Learn can be supported with this very cost-effective infrastructure.

The benchmarks were run using a single Dell PowerEdge R900 server for the database to show the excellent scalability of this cost-effective system. For customers that require a high availability option, additional database servers can be deployed in an Oracle RAC configuration and scalability is expected to be just as good, but has not been tested.

The following chapter provides additional details on configuring and sizing the reference architecture to meet an institution's specific needs.

Chapter 4

Best Practices and Configuration Guidelines

The benchmark tests conducted in Dell's labs provided a good baseline for defining the conditions that result in optimal performance for Blackboard Learn. Overall performance and management efficiency can also be impacted significantly by how the system is configured. The guidelines in this chapter provide best practices for configuring Blackboard Learn software to run well on Dell servers and storage solutions with Citrix XenServer virtualization software.

Why a Virtual Environment Makes Sense

It is common for customers to ask why there is so much focus on virtualization these days and why use multiple VMs with larger servers rather than simply using a lot of smaller servers. The main reason that virtualization makes sense with Blackboard Learn software is that the application can more fully utilize the CPUs in a virtualized environment. In part, this is because today's CPUs have become so powerful that running a single Blackboard Learn application instance on a single CPU server without virtualization leaves the CPU underutilized. By virtualizing a larger server and running multiple application instances across each CPU core, the CPUs are more fully utilized. A second point is that it is more cost-effective to add extra memory and CPU resources to a server so that it supports multiple VMs versus buying a complete server box.

Another benefit of the virtualized environment is that it offers an efficient means to process many small requests, which are typical in the Blackboard Learn application. This is as opposed to scientific applications such as fluid dynamics where a single algorithm can require a vast amount of CPU resources for a large calculation in which sequential processing is critical.

Application Tier Configuration and Sizing Guidelines

The performance capacity illustrated in Chapter 3 can be roughly translated into 7,500 active sessions per VM in the application tier. Considering the 12 VM configuration discussed in Chapter 3, the multiplier of 7,500 active sessions per VM would yield a total theoretical capacity of 90,000 active sessions for that configuration. The fact that the actual measured capacity was 102,590 active sessions is simply because the multiplier was rounded down to 7,500 to be conservative.

To reach its capacity of 7,500 active sessions, a VM should be configured as follows:

- 8GB of memory is needed to support a 4GB Java Virtual Machine (JVM) along with 2GB of memory for the Apache Web server and 2GB for the operating system and monitoring tools. See the section below on configuring the JVM for more information about why a 4GB heap is recommended for the JVM.

- A single VM requires two Intel Xeon 5500 series processor cores, which results in four virtual CPUs with Citrix XenServer. Four virtual CPUs showed the best performance with the virtualized environment with the Intel Nehalem processors. If a different processor is used, the configuration will need to be adjusted to account for the change in CPU capacity. Using an older or slower CPU, for example, would likely mean that more processor cores and more virtual CPUs would be required for each VM.
- The Citrix XenServer hypervisor requires an additional 2GB of memory above and beyond what is allocated for each VM. Thus if there are four VMs, a total of 34GB of memory is required (8GB * 4 + 2GB).

As mentioned earlier, performance will vary based on workload requirements. So, for workloads that are more demanding than that of the test scenario described in Chapter 3, a lesser multiplier than 7,500 should be used for the capacity of each VM.

Configuring the JVM

The Java object heap (Java heap) stores Java objects, including instances of classes and the data that the objects contain, such as primitives and references. The amount of memory allocated to the JVM as a whole and to its individual regions is different for every Blackboard Learn deployment. The appropriate size of the Java heap will vary depending on application usage.

With 64-bit operating environments, the Sun Java HotSpot Virtual Machine is no longer limited to the 1.7GB maximum heap size for 32-bit deployments. Testing by the benchmark team showed that a JVM heap size of 4GB is a good sweet spot. The 4GB heap size provided better performance and better availability for Blackboard Learn compared to smaller heap sizes. Larger heap sizes such as 8GB and 16GB were evaluated, but Blackboard’s best practices for availability and scalability recommend a more manageable JVM size.

Actual memory consumption can be monitored using operating system tools, or third-party tools such as Quest Foglight. If the memory resources are saturated or fully consumed, some adjustments will be necessary. The first recommendation is to allocate more memory to the JVM heap. If this does not work, consider adding additional JVMs to the deployment.

Using Non-standard (-XX) JVM Arguments

The Sun Java Virtual Machine contains non-standard HotSpot VM options that are used to improve virtual machine behavior, performance, and debugging. Non-standard options are prefixed with –XX.

Many non-standard options are available that can provide greater throughput and reduce overall latency. Tuning exercises can help in determining which options are best suited for the conditions in a specific customer’s production environment. Different logging options can be used to provide insight into the value of each option.

Table 1 describes the JVM options that will provide that majority of the performance gain.

JVM Option	Description
-XX:-UseConcMarkSweepGC	The concurrent collector is used to collect the tenured generation and does most of the collection concurrently with the execution of the application. The application is paused for short periods during the collection.
-XX:+UseParNewGC	This is used in conjunction with the concurrent collector on multiprocessor machines to enable young generation collection.
-XX:ParallelGCThreads	This option should be set to the number of available CPUs for concurrent garbage collection. This option is to be used in conjunction with -XX:+UseParNewGC.
-XX:NewRatio	This option sets the size ratio between young and old generational spaces. This is typically set to a value between two and four depending on garbage collection patterns.

Table 1: Non-standard JVM Options that can help improve performance

For more information about the JVM options supported by Java, visit java.sun.com/javase/technologies/hotspot/vmoptions.jsp. Given the variety of options and the chance that these options may change from release to release, Blackboard recommends that each customer test the impact of the different options within their environment by reviewing the output in the garbage collections logs.

Summary of Recommended Configuration Guidelines for the Application Tier

Table 2 summarizes the key recommended guidelines for configuring VMs in the application tier.

	Recommended Configuration
CPUs	2 cores per VM (Four threads/VM with hyperthreading)
Memory	8GB/VM (plus an additional 2GB for the hypervisor)
Operating System	Red Hat Enterprise Linux 5.3 64-bit edition
Hypervisor	Citrix XenServer 5.5 Dell Edition

Table 2: Configuring servers in the application tier

Additional information about configuring the Blackboard Learn application for maximum performance is available in the Blackboard white paper entitled, "Optimizing Blackboard Learn Deployments." This white paper can be obtained from a Blackboard representative.

Database Tier Configuration and Sizing Guidelines

The database must be sized to meet the requirements for the maximum transaction volume that can be generated by the application tier. The more VMs that are deployed in the application tier, the greater the potential load on the database. The number of Oracle processes active in the database server at a given point in time is dependent on the number of connections coming from the application servers. When more Oracle processes are active, more CPU and memory resources are needed to maintain high service levels in the database server.

As discussed in Chapter 3, the Dell PowerEdge R900 server configured with 48GB of memory and six processor cores demonstrated the ability to support 102,590 active sessions in the benchmark tests. From this performance information and an analysis of the CPU and memory resource utilization during other tests, the following configuration guidelines were developed for the database tier:

- One CPU core is needed for every two VMs in the application tier. This number was arrived at by simply considering that six cores were able to support the throughput of 12 VMs in the application tier. Testing also verified that a two-core database server could indeed support an application tier with 4 VMs.
- Memory requirements can be calculated by starting with 4GB to 8GB of memory for the System Global Area (SGA) and 2GB to 4GB for the OS and monitoring. An additional 2.4GB should be added for each VM in the application tier. This memory will support additional Oracle processes in the Program Global Area (PGA). The number of 2.4GB per VM was arrived at by noting that a single Oracle process requires approximately 3MB of memory and that a single VM can require as many as 800 Oracle processes during peak loads (3MB * 800 = 2.4GB).

As mentioned earlier, an Oracle RAC configuration may also be utilized in the database tier. In this case, additional database servers in the Oracle RAC configuration should be configured to match the requirements outlined herein for a single database server.

Table 3 summarizes the key recommended guidelines for configuring the database server.

	Database Configuration/Multiplier
CPU	0.5 cores/VM (One thread/VM with hyperthreading)
Memory	4GB to 8GB for the SGA, 2GB to 4GB for the OS and DB monitoring tools, and 2.4GB/VM for Process/PGA
Operating System	Red Hat Enterprise Linux 5.3 64-bit edition

Table 3: Configuration guidelines for the database server

Oracle Database Optimization

The configuration of the Oracle database and its storage infrastructure can significantly affect system performance for medium and large Blackboard Learn installations. The following guidelines may be helpful in optimizing Oracle database performance:

- The Oracle database should be allowed to allocate a memory space of larger than 3GB. This can be accomplished by applying the appropriate Oracle patches. Oracle Patch 7272646 is the appropriate patch for Oracle 11g v11.1.7.0.
- The Oracle database should be striped across a number of high-performance drives. The number of drives required depends on the expected volume of database transactions. RAID 10 configuration is recommended for the drives.
- Database performance can be optimized by using separate drives for different types of data stored in the database. There should be separate drives for the data files that comprise the database itself. Log files and tempdb files should be stored on other drives to reduce contention, allowing these files to be read or written in parallel with database reads and writes. This can help reduce I/O wait time.
- Blackboard customers may find the need to incorporate new indexes or change the behavior of the Oracle optimizer in order to achieve the best overall performance. The following chapter explains how Quest Foglight and Foglight Performance Analysis for Oracle can be used to help identify database performance issues to enable Blackboard customers to most effectively utilize the Oracle database.

For institutions seeking to dig more deeply into Oracle performance optimization, general information about Oracle optimization is abundantly available on the Internet. There is also a detailed section called, "Optimizing Oracle Database," in the Blackboard white paper entitled, "Optimizing Blackboard Learn Deployments." This white paper can be obtained from a Blackboard representative.

Network Configuration Guidelines

The network environment is designed to spread the user load evenly across all VMs in the application tier and to maximize I/O throughput to the storage tier. Because there is a large amount of I/O in the Dell Reference Architecture for Blackboard Learn, it is important to configure the network properly to support transfer of a significant number of large files.

The recommended guidelines for the network configuration include:

- Servers in the application tier should have at least two NIC ports dedicated to iSCSI traffic.
- The database server should have four NIC ports dedicated to iSCSI traffic.
- Dell PowerVault NX1950 unified network storage system should have four NIC ports dedicated to iSCSI traffic and two NIC ports dedicated to NFS file server traffic.
- Extreme Summit switches used for iSCSI traffic should have “jumbo frames” enabled and “spanning tree portfast” enabled.
- The Citrix NetScaler application delivery controller should be configured to turn on TCP Offloading, Client Keep Alive Session, CookieInsert, Compression, and Integrated Caching. These settings were shown to significantly improve transaction response times during testing of the reference architecture.

Chapter 5

Maintaining Performance and Availability

Maintaining performance and availability of the Blackboard environment without assigning a big IT staff requires having the right tools for efficient monitoring and management. The Quest systems management tools were instrumental in helping the project team more than double the performance of the Dell Reference Architecture for Blackboard Learn during the benchmark tests and can offer similar benefits to Blackboard customers. The tools enabled the benchmark engineers to peek under the hood of the Blackboard application layer and look within the Java Virtual Machine to see JVM memory utilization and garbage collection patterns. Foglight also helped identify which pages (transactions) were really being affected with increased loads and where the time was being spent within the application stack, thus enabling the team to direct their optimization efforts with precision.

Blackboard customers can use Quest's Application Management Suite to help optimize performance and to pinpoint performance bottlenecks in order to maintain high service levels over time. Quest Software provides all of the tools needed to simplify administration and optimize availability and performance.

More Efficient Application and Database Management

Academic institutions often have limited resources for system administration and management, making it a challenge to maintain high performance for key infrastructure components such as the Tomcat application server and the Oracle database used by Blackboard Academic Suite. Quest's Foglight® provides a view into the entire Blackboard application from end user to database for a complete understanding of the entire infrastructure. This helps reduce or eliminate service disruptions and can make it easier to meet application availability and performance requirements.

Foglight® not only monitors the Blackboard infrastructure (applications, databases, web servers, network and virtual servers), but it also monitors end-user performance to track usage levels, usage patterns and session analysis down to the individual user.

Quest Software provides the following tools to enable comprehensive performance management of the Blackboard application, the Oracle database and the virtualized environment of the application tier.

Foglight®

Foglight® is an enterprise-class solution providing true 24x7 monitoring, custom threshold alerts, advanced workload analysis, and automated optimization. It can be used to monitor the end-to-end Blackboard Academic Suite application down to individual Java virtual machines, including exposing the relationships and interaction between all components in the application environment. Foglight provides application- and Oracle-specific views and alerts, enabling a faster, more intuitive path to the problem and then to the solution.

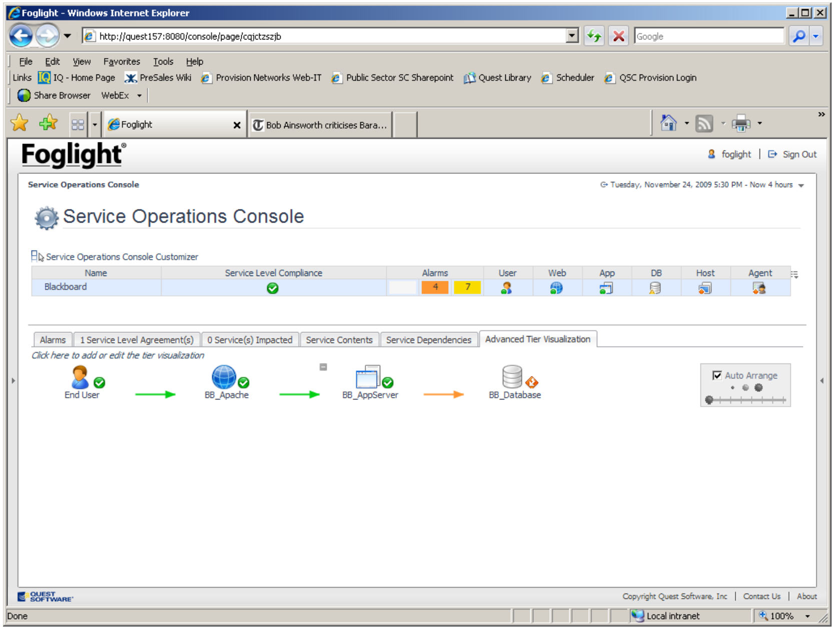

Figure 4: Foglight provides true 24x7 monitoring.

Foglight® Database Performance Analysis for Oracle

Foglight® Database Performance Analysis for Oracle provides detailed historical and real-time analysis of performance issues. Its performance management digital dashboard can help DBAs efficiently identify performance issues and quickly isolate the root cause. Using Quest's proprietary StealthCollect™ technology, Foglight Database Performance Analysis enables administrators to quickly correlate application SQL performance with system resource consumption. The broad range of advanced diagnostic features in Foglight Database Performance Analysis gives DBAs comprehensive visibility as well as detailed granular views of performance issues.

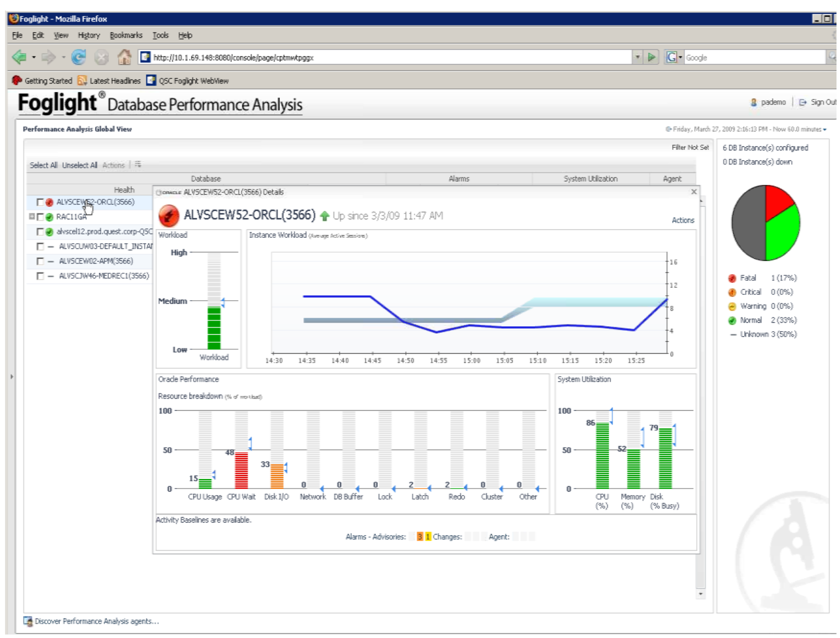

Figure 5: A performance management digital dashboard

Spotlight® on Oracle

Spotlight® on Oracle is an easy-to-use performance monitoring solution that can help administrators quickly and efficiently discover, diagnose, and resolve performance problems within an Oracle database environment. Deep diagnostics enable quick identification of the root cause of a performance bottleneck. SQL performance can be optimized through detection of problematic SQL, automated SQL tuning and benchmarking capabilities. Real-time alerts also allow administrators to respond more quickly to performance problems.

Figure 6: Spotlight on Oracle helps administrators quickly identify root causes.

Toad® for Oracle

Toad® for Oracle can help Oracle database professionals be much more efficient with tasks such as writing or tuning Oracle SQL queries, generating reports, and maintaining overall database performance. It provides the power to not only help DBAs build great code and prevent costly coding errors, but can also simplify management and maintenance of Oracle databases. As the industry's leading, proven Oracle database development and administration tool, Toad for Oracle can help ease the IT resource constraints common to many academic institutions by making IT staff more productive.

LiteSpeed® Engine for Oracle

LiteSpeed® Engine for Oracle offers administrators an advanced tool that can help cut storage costs and reduce backup and restore times by up to 50 percent. Integrating seamlessly into Oracle's Recovery Manager (RMAN) and Export backup strategies, Oracle DBAs charged with handling secure backup and recovery strategies on a lean budget can realize 70 to 90 percent compression with LiteSpeed while maintaining complete control over the backup and recovery process.

Chapter 6

Summary

Together, Dell, Blackboard, Citrix, and Quest have created a scalable reference architecture that allows customers to achieve responsiveness, scalability, and predictability of performance with a cost-effective infrastructure. The virtualized infrastructure of Dell servers and storage systems utilizing Citrix XenServer provides flexibility, scalability, and very high system utilization. The Oracle database has been optimized for higher performance and takes advantage of high performance iSCSI storage for higher throughput.

The end-to-end management infrastructure based on monitoring and management tools from Quest Software can help customers improve service levels as well as management efficiency. It provides administrators with effective tools that can simplify operational management and help improve the performance of the virtualized infrastructure on which the reference architecture is based.

Institutions that utilize the reference architecture can benefit from:

- **Higher service levels** – The architecture is designed to help optimize service levels and provide enhanced reliability through mature operation of the datacenter.
- **Reduced cost** – Citrix XenServer virtualization technologies enable a consolidated solution with higher resource utilization and the combined solution offers incredible performance on a cost-effective platform. Best practices for properly configuring and managing the IT infrastructure can also reduce the cost of deploying and maintaining the online learning environment.
- **Faster time to delivery** – The reference architecture saves institutions time by greatly reducing the need for research or trial and error discovery when building and optimizing an online learning environment.
- **Reduced risk** – Validated hardware configurations and software partners are combined with best practices for operational performance to greatly reduce the risk of unforeseen problems in a production implementation of Blackboard Learn.
- **Simplified management** – Tools and best practices for maintaining high service levels can save time and simplify the management process.

For More Information

The Web links in Table 4 provide resources for additional information about the technologies and solutions from Dell, Blackboard, Citrix, and Quest that comprise this reference architecture.

URL	Description
www.dell.com/hied	Dell solutions for higher education
www.dell.com/poweredge	Dell PowerEdge servers
www.dell.com/equallogic	Dell EqualLogic PS5000 Series iSCSI SAN array
www.blackboard.com	Blackboard Inc. home page
www.quest.com/oracle	Quest Management tools for Oracle databases
www.citrix.com/xenserver	Citrix XenServer server virtualization platform
www.citrix.com/netscaler	Citrix NetScaler application delivery solutions
www.citrix.com/site/resources/dynamic/partnerDocs/Citrix-XenServer-Equallogic_final.pdf	White paper entitled, “Deploying Citrix XenServer 5.0 with Dell EqualLogic PS Series storage”

Table 4: Web Links for Additional Information

Appendix A

Components of the Reference Architecture

Tables 5 and 6 describe the recommended hardware and software components within the reference architecture.

Hardware Component	Description
Dell PowerEdge R900 server	The PowerEdge R900 server is Dell's most powerful Intel server, offering up to 24 Intel Xeon processor cores. It is designed and optimized to provide outstanding virtualization performance.
Dell PowerEdge R710 server	Dell's new 2-socket 2U rack server is ideal for business-critical applications such as Blackboard Learn. The PowerEdge R710 server utilizes the latest Intel Xeon X5500 series processors and supports up to 125% more memory capacity and more integrated I/O than the previous generation PowerEdge 2950 III. This increased capacity is crucial for virtualization performance and scalability.
Dell EqualLogic PS5000E iSCSI Array	The PS5000E array is an intelligent virtualized iSCSI storage array that is designed to provide rapid installation, simple management, and seamless expansion. Its SATA disk drives provide cost-effective high capacity storage for the application tier.
Dell EqualLogic PS5000XV iSCSI Array	The Dell EqualLogic PS5000XV array is a virtualized iSCSI SAN solution that combines intelligence and automation with fault tolerance to provide simplified administration, rapid deployment, enterprise performance and reliability, and seamless scalability. With 15,000 RPM SAS disk drives, the PS5000XV array provides high transactional performance, making it an ideal SAN platform for the database server in this reference architecture.
Dell PowerVault NX1950 unified network storage solution	The Dell PowerVault NX1950 is a unified network storage solution that simultaneously stores both file and application data. It can be used as an integrated storage solution or a NAS gateway and is deployed in the reference architecture as a NAS gateway connecting to the Dell PS5000E storage system for application tier storage.

Hardware Component	Description
Citrix NetScaler Web application delivery controller	The Citrix NetScaler Web application delivery controller functions as an application accelerator through caching and HTTP compression. It provides advanced traffic management through load balancing and content switching functions for Layers 4-7 of the network stack.
Extreme Networks BlackDiamond 8810 switch	BlackDiamond 8800 series switches provide the ideal core network for a medium-sized network with high-performance and high density 10 Gigabit Ethernet and Gigabit Ethernet interfaces. Blackboard customers can connect up to 192 ten-gigabit ports or 768 gigabit ports in a single 14RU BlackDiamond 8810 system.
Extreme Networks Summit X450a-48t switch	Based on the revolutionary ExtremeXOS operating system, the highly flexible and scalable Summit X450a series of switches are ideal for traditional small core enterprise networks and are excellent for first level aggregation in a medium or large network. Its compact 1RU format with gigabit and optional 10 Gigabit Ethernet ports provides great versatility and offers a full range of Layer 2 to Layer 4 functionalities on every port to allow high productivity.

Table 5: Reference Architecture Hardware Components

Software Component	Description	Version
Blackboard Learn software	All Blackboard license holders install the same Blackboard Learn software. The available capabilities are controlled by the institution license. In addition to Blackboard Learn course delivery capabilities, the following capabilities are available to be licensed: community engagement, content management, and outcomes assessment.	Release 9
Oracle 11g database	Oracle Database 11g delivers industry leading performance, scalability, security and reliability.	Release 1 Enterprise Edition
Citrix XenServer Dell Edition	XenServer is a server virtualization platform that offers near bare-metal virtualization performance for server and client operating systems. Citrix XenServer Dell Edition can be pre-installed on Dell PowerEdge servers, enabling rapid deployment of the virtualized Blackboard Learn environment.	Version 5.5

Software Component	Description	Version
Quest Foglight	Foglight is a complete performance management solution for enterprise database environments, addressing all aspects of database performance from 24x7 monitoring to advanced analytics and automated tuning.	5.5
Quest Foglight Performance Analysis for Oracle	Foglight Performance Analysis for Oracle provides a consolidated, integrated view to perform deep diagnostics and intelligent analysis, track changes, and enable root-cause detection of performance concerns in Oracle database environments. It captures and analyzes database workload information to identify performance issues and isolate their ultimate root cause—whether it's a specific user, program, action or machine.	6.5.4
Quest Toad for Oracle	Toad for Oracle is an award-winning solution that gives DBAs and IT administrators the productivity tools needed to proactively manage the Oracle database, help optimize performance, and maintain high service levels.	10.0
Quest Spotlight on Oracle	Spotlight on Oracle provides real-time diagnostics to help eliminate bottlenecks in Oracle and Oracle RAC environments.	7.0
Quest LiteSpeed Engine for Oracle	LiteSpeed Engine for Oracle provides fast, flexible backup and recovery with industry-leading compression technology.	1.2
Red Hat Enterprise Linux	Red Hat Enterprise Linux 5 is the world's leading open source application platform.	V5.3, Patch Update 3 for 64-bit

Table 6: Reference Architecture Software Components

Appendix B

Workload Activities in the Test Scenario

Table 7 shows the breakdown of activities that comprise the test scenario used for benchmarking. These activities are representative of the user activity patterns found in Blackboard customers who have a large population of users who are using the application for distance learning.

Of all of these different types of transactions in the test, 15% represented the use of features that were new in release 9 of Blackboard Learn and 85% utilized features that already existed in the previous release. The test case scenario also consists of 91% student-related activities and 9% teacher activities.

Transaction Type	Percentage of Total Transactions
Creating and reading announcements	5%
Creating, reading, and replying to messages on the Discussion Board	20%
Taking and grading quizzes and exams	10%
Creating and reading entries from Group Blogs and adding comments	10%
Submitting and grading assignments	10%
Downloading and opening course content including PDFs, Microsoft Word documents, etc	20%
Viewing external links	5%
Creating and reading entries and adding comments in the Journal	5%
Reviewing and scoring grades and creating grade reports in the Grade Center and Performance Dashboard	15%

Table 7: Workload activities that comprise the test scenario