

ESG Lab Review

Vblock Specialized Systems for High Performance Databases

Date: December 2013 **Author:** Aviv Kaufmann, ESG Lab Analyst and Mike Leone, ESG Lab Analyst

Abstract: This ESG Lab review documents hands-on testing of the Vblock Specialized Systems for High Performance Databases focusing on high performance in an Oracle 11g Release 2 environment and system reliability in the event of component failures.


The Challenges

Databases only perform as well as the systems they run on. Designing a system to power a true high performance database requires design choices that offer extreme performance, interoperability, and specific knowledge of tuning parameters that help streamline the operation of the database. The engineering resources and expertise required to design such a system are not generally available to all organizations. This drives many organizations to turn to specialized consulting services to help them piece together a high performing database system using high-end specialized components from various hardware vendors.

Integrated computing platforms (ICPs) are platforms in which servers, storage, network connectivity, and (in some cases) software are combined in a single solution. ICPs have proven to save businesses countless hours designing, configuring, and testing servers, networks, and storage subsystems. ESG research respondents indicated that ease of management, faster deployment times, and improved total cost of ownership (TCO) are some of the benefits that ICPs offer their organizations (see Figure 1).¹

Figure 1. Top Six Benefits of Integrated Computing Platforms (ICPs)

What benefits do you believe an integrated computing platform offers your organization? (Percent of respondents, N=471, multiple responses accepted)



Source: Enterprise Strategy Group, 2011.

¹ Source: ESG Research Brief, [Integrated Computing Trends](#), March 2011.

The Solution: Vblock Specialized Systems for High Performance Databases

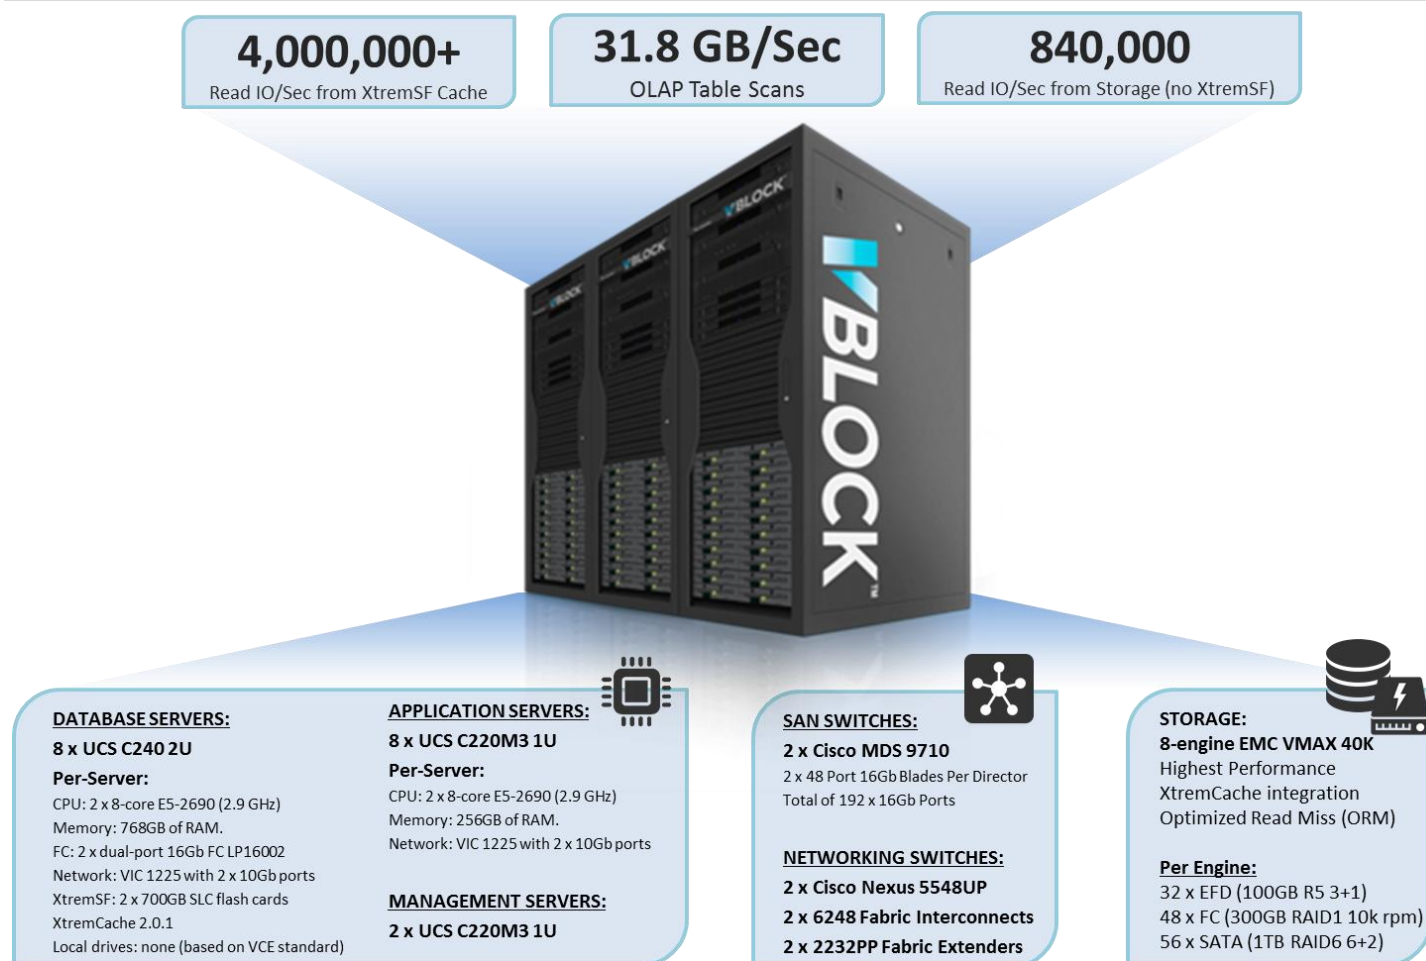
The [VCE](#) Vblock System is a well-known and extensively deployed converged infrastructure platform that combines best-of-breed technologies from VMware, Cisco, and EMC into a powerful integrated computing platform. The new Vblock Specialized Systems for High Performance Databases are designed as a simple solution to the most demanding transactional and response-time-sensitive databases required by today's businesses.

Like all Vblock Systems, the Vblock Specialized Systems for High Performance Databases are comprised of compute (database, application, and management servers from Cisco), network (cutting-edge, three-layer network and SAN technologies from Cisco), and storage (EMC's VNX or VMAX systems and XtremSF cards). Although this lab validation was completed on a bare-metal configuration, VMware vSphere may also be used for virtualized database configurations. The Vblock Specialized Systems for High Performance Databases are specifically designed to deliver the lowest possible response times and highest possible throughput by:

- Maximizing the amount of memory and processing power in the database servers.
- Utilizing dual PCIe flash, host-based XtremSF cards along with XtremCache software.
- Combining the low-latency benefits of 10GbE and 16Gb FC technologies with Director Class Switches.
- Offering the ideal mix of flash, performance, and capacity drives optimized with FAST-VP autotiering.

The configuration details of the system under test are shown in Figure 2.

Figure 2. Vblock Specialized Systems for High Performance Databases



Getting Started

ESG Lab began testing by touring the cutting edge VCE facility in which new Vblock Specialized Systems for High Performance Databases are built, pre-configured, and pre-tested prior to being shipped to the customer. The facility was laid out in a highly efficient manner that made use of patented processes to speed production, insure quality, eliminate configuration errors, and label/document all system connections. ESG Lab inspected the physical components of the largest system offered by VCE, which included seven racks containing all the storage, server, and networking resources. The entire process was designed to save customers months of work and bring their databases to production as quickly as possible. In some cases, specialized software installations can occur onsite. It's important to note that though ESG Lab testing focused on Oracle RAC, the Vblock Specialized Systems for High Performance Databases are built to take advantage of other large, enterprise databases as well.

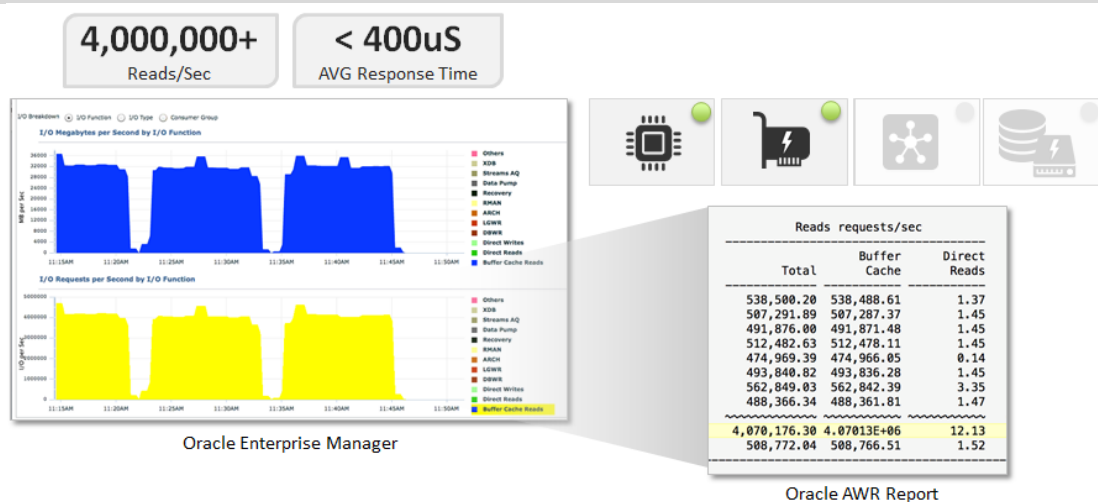
Performance Testing

ESG Lab validated the performance of the Vblock Specialized Systems for High Performance Databases running an eight-node Real Application Clusters (RAC) configuration of an Oracle Database 11g Release 2 on RHEL 6.3. All performance testing was completed on a simulated real-world populated database. Testing focused on measuring the throughput and latency of Oracle SQL-driven I/O and the bandwidth driven during table scans. ESG Lab utilized the widely adopted and publicly available Silly Little Oracle Benchmark kit ([SLOB](#)) to efficiently generate realistic system-wide, random, single block, application-independent SQL queries. The SLOB benchmark exercised all components of the Vblock Specialized Systems for High Performance Databases by stressing the physical I/O layer of Oracle through SGA-buffered random I/O without being limited to a specific load-generating application.

First, ESG Lab measured the maximum system performance achieved when utilizing a data set that fit entirely into the host-based XtremCache. Each database server contained two 700GB SLC XtremSF flash cards and XtremCache 2.0.1. The cards are used as a write-through cache to retain a copy of the data and cache read I/O.

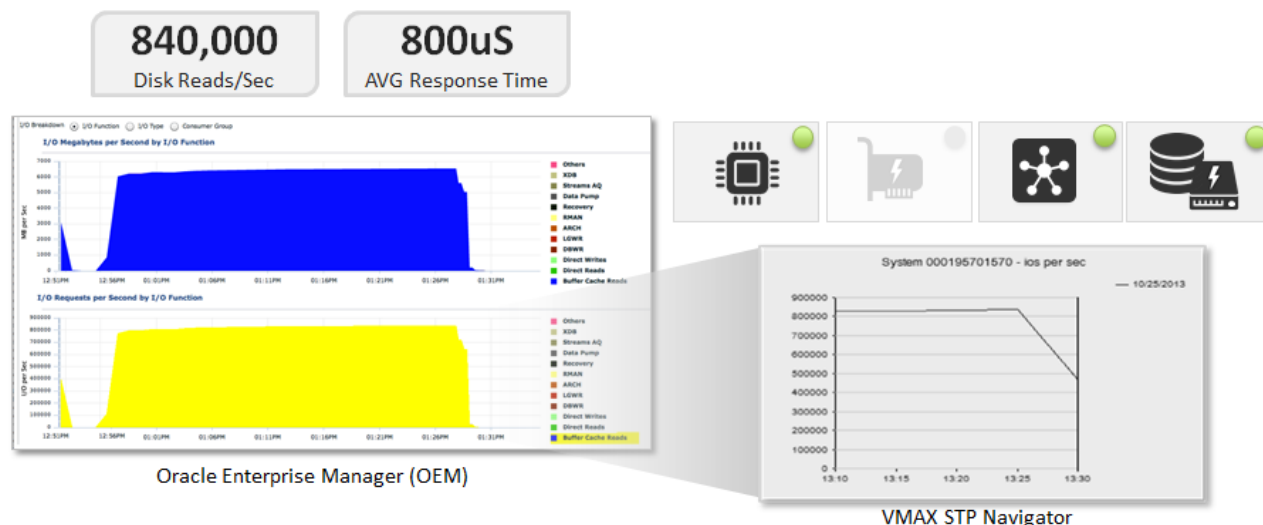
The SLOB workload was configured to simulate 576 users accessing a 1.5TB active data set. The high transactional workload generated a heavy CPU burden on the database servers. VCE designed the system under test with minimal tuning parameters to best simulate a default Oracle initialization. ESG Lab witnessed high, sustainable performance of over four million random Oracle database I/Os per sec while delivering a very low 400 microseconds response time. Highlights of the results are shown in Figure 3.

Figure 3. OLTP Performance with XtremSF Host-based Cache Cards



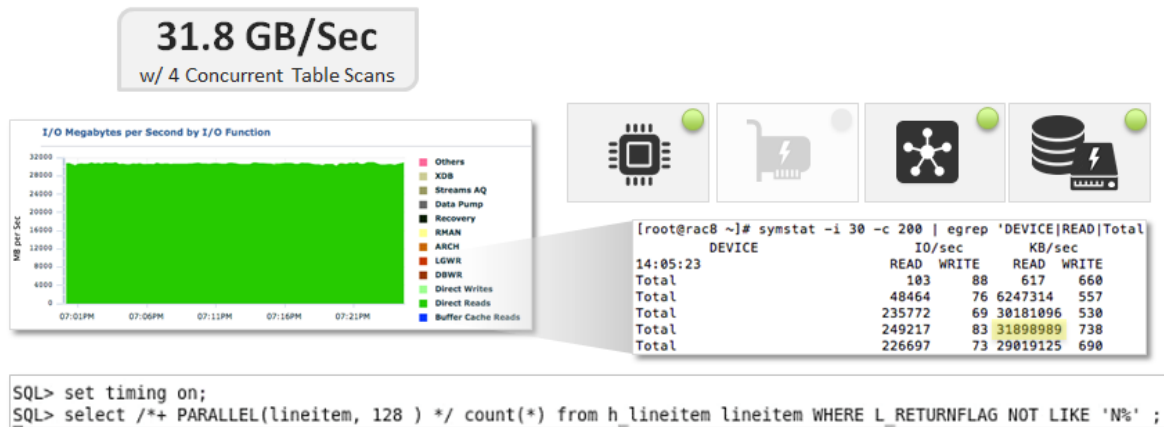
Next, ESG Lab validated the performance of the system while performing random Oracle database physical I/O to the VMAX 40K Storage system. The XtremCache was disabled prior to testing to simulate a scenario in which transactions overflowed from an active XtremSF and XtremeCache environment. The number of concurrent users used to generate the SLOB workload was increased to 1,024 to allow for a greater degree of parallelism. ESG Lab verified that the storage system was able to deliver 840K disk reads/sec at a very low average response time of 800 microseconds using both Oracle Enterprise Manager and VMAX STP Navigator. The test results are shown in Figure 4.

Figure 4. OLTP Performance with VMAX 40K Storage Subsystem



Finally, ESG Lab validated that the Vblock Specialized Systems for High Performance Databases were able to provide high throughput performance when running table scans. ESG Lab generated four concurrent SQL queries that scanned a 2.85TB line item table and monitored the performance using both Oracle Enterprise Manager and symstat. Once the test run completed, the throughput results were validated by auditing the Oracle AWR report, where ESG Lab verified very high throughput performance of nearly 32 GB/sec. The results of the OLAP performance testing from Oracle Enterprise Manager and symstat are shown in Figure 5.

Figure 5. Measured Throughput Performance with VMAX Storage System



Why This Matters

Modern performance-oriented database workloads require high throughput, low latencies, and flexible scalability to keep up with today's demands for data. Spending less time accessing data can mean hundreds of thousands of dollars of revenue to an organization. Designing a database with the ability to scale and deliver high performance for mixed workloads like OLTP and table scans within the same system is a challenge for even the most skilled database architects.

ESG Lab validated that a single configuration on the Vblock Specialized Systems for High Performance Databases performed over four million I/Os per second with a 400 microseconds average response time for OLTP workloads. The system also delivered close to 32GB/sec while performing table scans. Designing a database system to achieve these high levels of performance would be difficult for any sized organization, but VCE made the configuration and delivery process quick and easy.

Availability and Reliability Testing

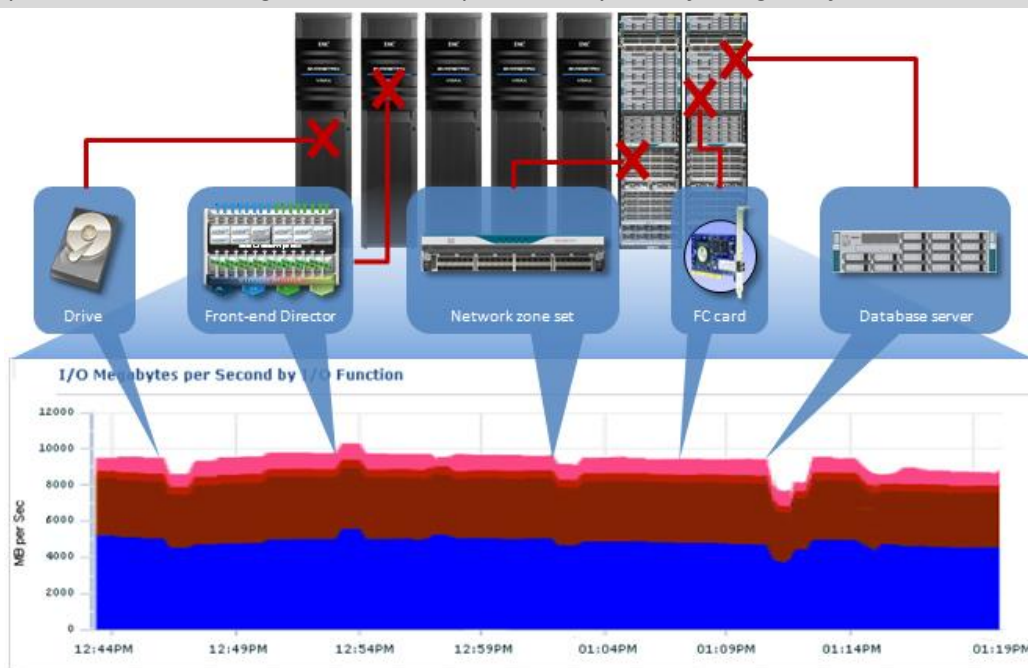
Although the Vblock Specialized Systems for High Performance Databases are designed to deliver extremely high levels of performance, every component in the system is also designed to provide extremely high levels of availability. The Vblock Specialized Systems for High Performance Databases combine the high availability of Oracle RAC with the highest availability hardware components from Cisco and EMC. VMAX data protection features such as backup, recovery, and continuous availability give IT administrators the ability to meet the most demanding SLAs. ESG Lab validated the reliability of the Vblock Specialized Systems for High Performance Databases by testing the system's ability to withstand component failures and use storage-level clones to quickly recover a corrupt database.

First, ESG Lab tested the ability of the system to maintain high levels of performance while dealing with multiple component failures. The SLOB benchmark tool was used to generate a moderately heavy mixed OLTP workload to the system. Each host was configured to perform SQL SELECTs and UPDATEs with 256 concurrent sessions to a total of 640GB of data. VMAX tools verified that the workload consisted of about 50,000 read I/Os per sec and 231,000 write I/Os per sec to disk. Oracle Enterprise Manager was used to monitor the overall database performance of about 9,500 MB/sec.

Once the workload hit a steady state, ESG Lab systematically forced component failures at each level of the system. Resiliency of the VMAX storage system was tested by pulling out multiple disks followed by the powering off of an entire front-end director. The workload continued to run on all hosts and ESG Lab verified that the faults resulted in minimal interruption to the overall level of system performance using Oracle Enterprise Manager.

Next, ESG Lab simulated a failed SAN connection (switch, cable, or host FC card) by removing a zone set from the Cisco MDX switch configuration. The workload continued to run uninterrupted as the multipath drivers rerouted the workload and the Oracle RAC configuration simply swapped logs. Finally, ESG Lab simulated the failure of an entire database server in the cluster by removing all of the active network connections. The Oracle configuration was monitored to verify proper removal of the node from the RAC cluster. A slight drop in system performance was expected as one-eighth of the workload generators were removed, but because the systems were not fully saturated prior to the failure, performance remained relatively high. The workload was then redirected to one of the surviving nodes of the cluster, and system performance returned back to the original level of performance. Figure 6 highlights the results of the component failure testing. Each hardware blowout from the performance chart shows the time of that particular hardware failure and the impact the failure had on the overall system performance.

Figure 6. Component Failure Testing on the Vblock Specialized Systems for High Performance Databases

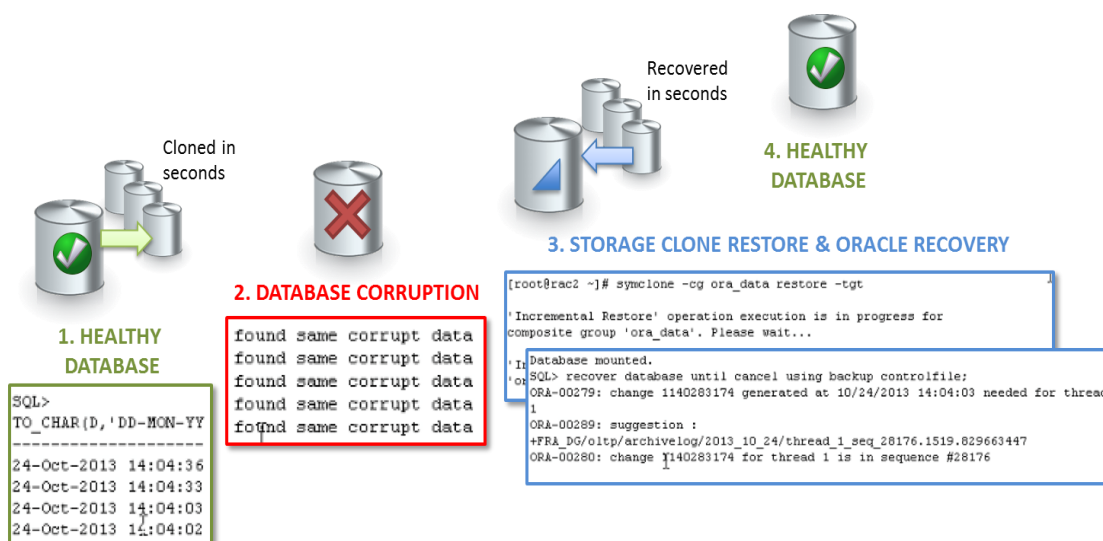


After the component failure testing, ESG Lab tested the ability of the system to quickly recover from a database corruption. A composite group was created on the VMAX storage system consisting of the Oracle OLTP database LUNs. A table was created in the database and populated with current timestamps. An Oracle Backup was created while ensuring DB and ASM metadata consistency and VMAX thin clones of the composite group were quickly created. ESG Lab then corrupted the database by overwriting the source LUNs, which contained key components of the Oracle database. This was verified by identifying a number of data corruptions in the Oracle logs.

ESG Lab stopped the database and ASM and verified that all instances were no longer running on any of the DB servers. Because the VMAX was aware of the data being stored on the XtremSF cache cards for each DB Server, the cards were first deactivated. ESG Lab then attempted a restore from the SymClone devices and in a matter of seconds, the restore completed.

Finally, an Oracle recovery was performed by overwriting the corrupted control files with the backup control files, mounting the database, and rolling forward the uncommitted transactions. Once the Oracle recovery completed, the database was restarted, the table containing the timestamps was successfully read, and the timestamps were verified to match the original data. The entire process, while appearing complex in nature, was completed using scripts in a matter of minutes. A high-level overview of the database corruption and recovery is shown in Figure 7.

Figure 7. Recovery of a Corrupt Database



Why This Matters

High availability is critical for all tier-1 databases. According to ESG research, 17% of organizations stated that any downtime whatsoever would result in significant revenue loss or business impact, while 74% of organizations stated that anything longer than one to three hours of downtime would result in a significant loss.² It is safe to assume that organizations that deploy a higher performance database certainly fall into the category of those that cannot tolerate any downtime.

ESG Lab validated that the Vblock Specialized Systems for High Performance Databases are capable of continuously maintaining a high level of performance in the event of simultaneous component failures at all levels of the system. An Oracle database was purposely corrupted, recovered, and placed back online with no loss of data. Using a combination of VMAX SymClones and Oracle backup and recovery procedures, ESG Lab completed the full database recovery in just minutes.

² Source: ESG Research Report, [2010 Data Protection Trends](#), June 2010.

The Bigger Truth

Performance and reliability are absolutely critical in today's high performance databases. Databases are constantly demanding higher levels of throughput, lower response times, and high levels of data protection, all while continuing to give the core business a competitive edge. ESG research shows that organizations with database applications such as OLTP, OLAP, and ERP are by far the leading adopters of high performance, low-latency, solid-state storage.³ But simply adding solid-state storage to your existing infrastructure will not solve the problem. Designing a high performance database that meets the demands of today's business can be a daunting task and few organizations have the expertise or the time to accomplish this in-house.

The Vblock Specialized Systems for High Performance Databases are designed to make it simpler and quicker for an organization to deploy a high performance database. The system is designed for fast and easy ordering using the highest performing and most reliable components from Cisco and EMC. In addition to fast servers, switches, and storage, highly available XtremSF host-based cards, along with XtremCache, are utilized to ensure the highest throughput and lowest response times. The system is pre-designed, pre-configured, pre-installed, pre-tuned, and pre-tested by expert database architects at VCE before delivery to the customer. Once it hits the customer floor, this product is ready for deployment—lowering the TCO by saving the customer months of internal work and increasing the ROI by speeding the time to deployment.

ESG Lab validated the high performance capabilities of eight-node Vblock Specialized Systems for High Performance Databases using an Oracle Database 11g Release 2 in a RAC configuration.⁴ The system performed over four million random physical I/Os per second driven by Oracle OLTP-like SQL queries at an impressively low 400 microseconds average response time, and delivered close to 32GB/sec while performing just four concurrent OLAP table scans. ESG Lab also witnessed the reliability and high availability of the solution by ensuring that the system could continue to operate and perform at a high level, even while multiple component failures were injected throughout various levels of the system design. This blend of extremely high performance and high reliability perfectly suit the needs of today's most demanding high performance databases.

By offering the Vblock Specialized Systems for High Performance Databases, VCE has taken the next step in integrated computing. VCE is not only offering the customer a best-of-breed, generalized platform in the Vblock 700, Vblock 300, Vblock 200, and Vblock 100 with which to build an infrastructure, but it is also helping to solve a complex and costly business problem for organizations. If your organization is preparing to spend months designing, testing, tuning, and troubleshooting a high performance database, ESG Lab strongly suggests letting VCE do the heavy lifting for you.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

³ Source: ESG Research Report, [Solid-state Storage Market Trends](#), November 2011.

⁴ Though this paper focused on the eight-node database server configuration, it should be noted that VCE also provides two- and four- node configurations where customers can not only run RAC- and non-RAC-based Oracle solutions, but also other enterprise-grade database applications.