Performance evaluation of Oracle 10*g* on HP StorageWorks Enterprise Virtual Array (EVA)

Proven performance and flexibility to meet mission-critical OLTP requirements

Table of contents

invent

Introduction

With over 200,000 customers, Oracle® databases are now managing exabytes (millions of terabytes) of data, which will only continue to grow exponentially each year. It is an enormous volume of information—much of it mission-critical, requiring fast, reliable access on a 24x7 basis. Therefore, the storage systems on which these data are stored must offer outstanding performance, availability and scalability.

HP StorageWorks Enterprise Virtual Array (EVA) is the ideal platform for Oracle databases. As a leader in the storage arena, HP has a keen understanding of the critical issues Oracle database administrators (DBAs) face today, including disk optimization, cloning of production volumes for testing and shrinking back-up windows. HP StorageWorks EVA addresses these issues by offering:

- Extensive scalability with very attractive price-performance
- Intelligent management of multiple arrays and hundreds of terabytes from a single, centralized location
- High throughput and zero downtime, with no single point of failure and rapid disaster recovery
- A compact footprint to conserve valuable floor space

One of the greatest concerns for Oracle DBAs is performance. HP and Oracle have conducted extensive testing using the HP StorageWorks Enterprise Virtual Array 5000 (EVA5000) with Oracle 10*g*, demonstrating that the HP StorageWorks EVA platform provides both performance and flexibility in supporting demanding transactional workloads. In this paper, HP and Oracle discuss how HP StorageWorks EVA offers the features most needed in Oracle database environments. We then provide details of the performance tests, including system configurations, test descriptions, results and analysis.

With its advanced virtualization and snapshot/cloning capabilities, the HP StorageWorks EVA storage platform stands out from the competition, enabling Oracle DBAs to increase the performance and flexibility of Oracle deployments. It is a vital component in the HP Adaptive Enterprise strategy in which business and IT are synchronized to capitalize on change—offering simplicity, agility and value over the long term.

Performance tests

HP conducted a series of tests, subjecting the HP StorageWorks EVA to various loads in an Oracle 10*g* environment. In this section, we provide hardware configurations, database layout, test descriptions and analysis for each of these tests.

System configurations

Server	Client	Storage
• (1) N-class server (HP 9000 rp7400 server – PA-RISC) – 550 MHz (8-CPU) • HP-UX 11iv1 • HP StorageWorks Secure Path (SP) 3.0D software • (6) 2 Gb Fibre Channel (FC) host bus adapters (HBAs) – 4 used for EVA • (1) 16-port FC SAN switch • 32 GB memory • 44 GB swap • Additional storage for storing test results, kits, etc. • (1) 1 Gb network card for client access • Oracle database release 10*g*R1 • (8) 50 GB LUNs	• (1) N-class server (HP 9000 rp7400 server – PA-RISC) – 550 MHz (8-CPU) • HP-UX 11iv1 • (2) 2 Gb FC HBAs • 32 GB memory • 44 GB swap • Additional storage for storing test results, backup, etc. • (1) 1 Gb network card for server access • Oracle database release 10*g*R1	• HP StorageWorks EVA5000 • Array firmware revision – 30.20 • (84) 36 GB, 10K rpm disks • (2) controllers – (2) host ports each • 1 GB cache per controller • (4) FC connections to the SAN switch

Secure Path for the HP StorageWorks EVA

Secure Path for the HP StorageWorks EVA is a server-based, multi-pathing software program that enhances the EVA disk arrays. Secure Path for the HP StorageWorks EVA presents a single virtual path to the host when multiple physical paths are available for the same Vdisk (logical device). Secure Path for the HP StorageWorks EVA monitors every physical path to the device and determines the status of available storage units and physical paths through internal diagnostics. Path or channel failures are reliably detected without inducing false or unnecessary failovers. In case of path failure, Secure Path software automatically redirects pending and subsequent I/O requests to an alternate path. Failover and failback actions are transparent and non-disruptive to applications. The Secure Path load-balancing feature can be enabled or disabled.

Test environment

Database layout and configuration	Workload
• Test environment • Database layout and configuration Workload • Stripe and mirror everything (SAME) methodology • (48) LVM RAW devices • (2) redo logs • (42) data files • Redo log stripe size – 128 KB • Data stripe size – 64 KB • Oracle block size – 2 KB to 16 KB • Database size – ~300 GB • Archiving turned off • Asynchronous I/O turned on • ~8 GB SGA	• OLTP-type I/O-intensive application* • ~60% read and ~40% write • I/O size varies, 2 KB (75%) – 16 KB (25%) • Workload intensity – Medium load – 50 batch users – Heavy load – 100 batch users • Max CPU utilization – Medium load – ~58–67% – Heavy load – ~87–96% • Mirror Write Cache (MWC) and Read Cache (RC) on by default on the EVA

* The EVA array I/O throughput based on our heavy workload was measured at 89% of the array performance. This clearly indicates that the EVA was stressed.

Figure 1: Single disk group without HP StorageWorks Secure Path software

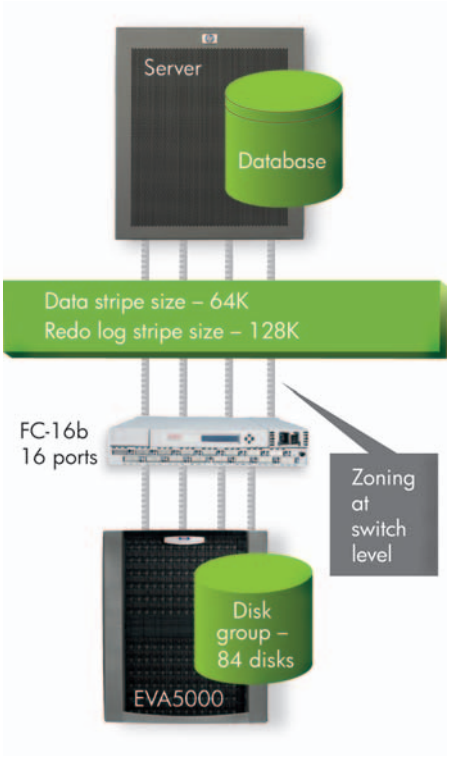

Table 1: Test results of EVA without Secure Path software, variation 1

	Logs—VRAID 1		Logs—VRAID 5		Logs—VRAID 0	
	Medium	Heavy	Medium	Heavy	Medium	Heavy
Transactional throughput/m	20,641	27,983	19,923	27,603	20,014	27,572
Log file parallel write (ms)	0-1	1-2	0-1	1-2	0-1	1-2
Db file sequential read (ms)	11	15	11	15	11	15
I/O throughput /sec	6,398	8,721	6,234	8,654	6,227	8,592

Test details

Oracle 10*g* was tested with HP StorageWorks EVA using a number of configuration variations, designed to evaluate performance for a range of potential scenarios.

The metrics used for these tests include:

- Transactional throughput – average number of transactions per minute measured during the run under stable condition
- Log file parallel write – latency of redo log file writes as measured using the Oracle "statspack" tool
- Db file sequential read – single block read as measured using the Oracle "statspack" tool
- I/O throughput – total number or physical reads and writes per second as measured using the Oracle "statspack" tool

In the following section, we describe each test, its results and our conclusions.

Test series 1 – single disk group

While Secure Path is required with the EVA, it is important to understand what impact the software may have on performance. This first set of tests was designed to determine the overhead of HP StorageWorks Secure Path software and how HP StorageWorks EVA performs with various features enabled or disabled. These tests were configured with VRAID 1 data files and redo logs on VRAID 1, VRAID 5 and VRAID 0. When the redo logs were configured on VRAID 0, Oracle mirroring (redo logs multiplexing) was used for data availability.

Single disk group without HP StorageWorks Secure Path software

The first step was to determine the performance of EVA without Secure Path software, as shown in Figure 1.

In this test, we first evaluated transactional throughput between logs on VRAID 1, VRAID 5 and VRAID 0 in a single disk group with zoning and no Secure Path. A zone is comprised of one EVA host port and one host bus adapter (HBA). With this configuration, we eliminated the multiple path masking normally provided by Secure Path when presenting any virtual disks to a single HBA. The results of this test are presented in Table 1.

As indicated by these results, the variation of transactional throughput between logs on VRAID 1, VRAID 5 and VRAID 0 is negligible.

Figure 2: Single disk group with HP StorageWorks Secure Path software

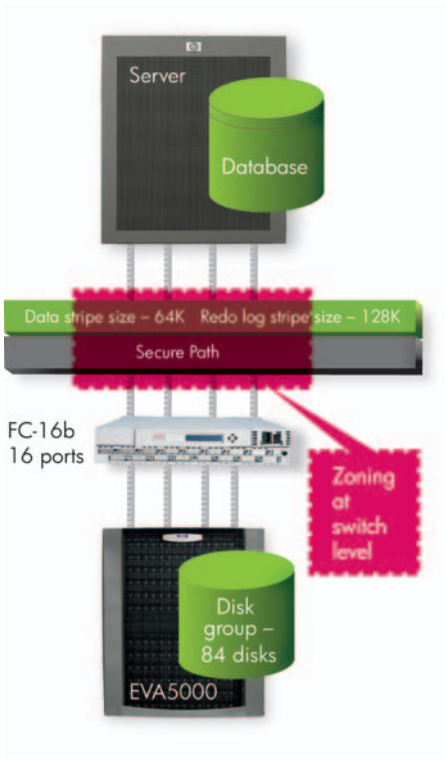

Table 2: Test results of EVA without Secure Path software, variation 1

	Read cache enabled		Read cache disabled	
	Medium	Heavy	Medium	Heavy
Transactional throughput/m	20,641	27,983	20,055	27,667
Log file parallel write (ms)	0-1	1-2	0-1	1-2
Db file sequential read (ms)	11	15	11	15
I/O throughput/sec	6,398	8,721	6,149	8,336

In a second variation of this configuration, we compared performance with EVA array read cache enabled and disabled. The results of this test are presented in Table 2.

Again, little variation in transactional throughput resulted from either having read cache enabled or disabled. Therefore, we recommend that read cache be kept on. Note that a very high Oracle read cache hit ratio (99–100 percent) mitigates changes in read cache settings.

Single disk group with HP StorageWorks Secure Path software

In the next set of tests, we included Secure Path software in the system configuration, again with zoning in a single disk group. (See Figure 2.)

The configuration includes VRAID 1 data files and redo logs with read cache enabled, and was tested in two scenarios: one with Secure Path software and one without. The results of this test are presented in Table 3.

Figure 3: Single disk group with Preferred Path feature of Secure Path with no zoning

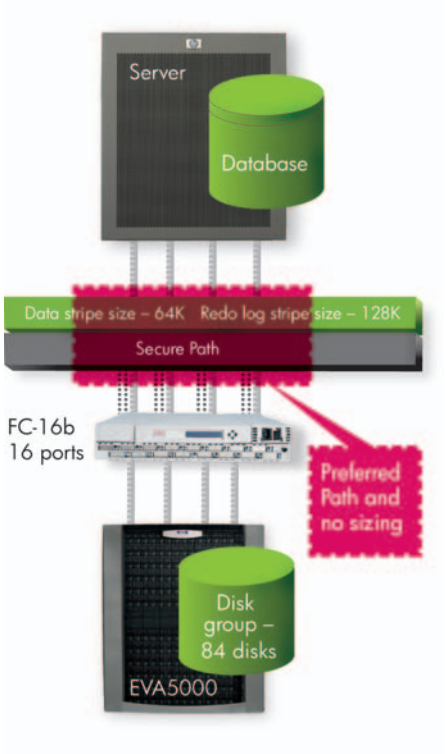

Table 3: Test results of EVA with and without Secure Path software

	With Secure Path		Without Secure Path	
	Medium	Heavy	Medium	Heavy
Transactional throughput/m	20,127	27,027	20,055	27,667
Log file parallel write (ms)	0-1	1-2	0-1	1-2
Db file sequential read (ms)	11	15	11	15
I/O throughput/sec	6,068	8,136	6,149	8,336

From the data collected in this test, we concluded that overhead from Secure Path is negligible and that CPU overhead is less than three percent.

In another variation of this configuration, we tested the effects of the Preferred Path feature of Secure Path with no zoning. (See Figure 3)

In the first scenario, we evaluated the Preferred Path feature for a given virtual disk (Vdisk). Vdisks can be accessed from any controller, but at any given moment there is only one active path to the Vdisk. When Secure Path is used, a single virtual path is visible to the host. Therefore, there is no need to use zoning unless there is a need to partition the SAN. The "spmgr" command displays the various physical paths to the Vdisks.

Figure 4: Single disk group without Preferred Path feature of Secure Path and no zoning

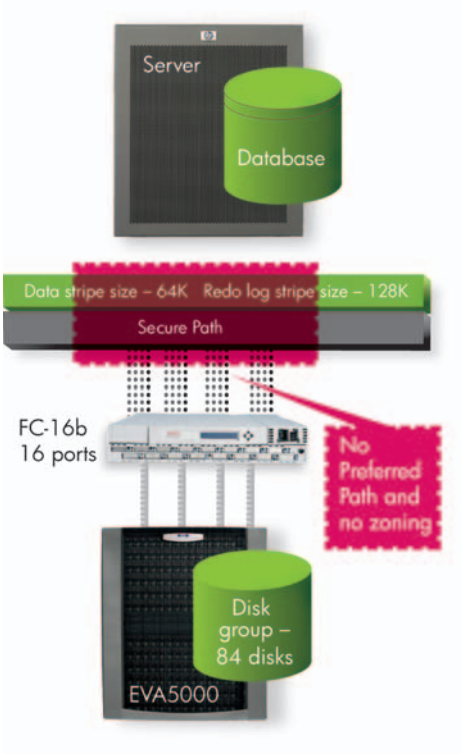

Table 4: Test results of EVA with and without Preferred Path feature of Secure Path

	With PP		Without PP	
	Medium	Heavy	Medium	Heavy
Transactional throughput/m	20,426	27,162	20,430	26,851
Log file parallel write (ms)	0-1	1-2	0-1	1-2
Db file sequential read (ms)	11	15	11	15
I/O throughput/sec	6,249	8,513	6,267	8,393

In the second scenario, we allowed Secure Path to direct I/O through a given controller without the use of the Preferred Path (PP) feature. (See Figure 4)

The results of these two scenarios are shown in Table 4.

From these test results, we determined that it is better to use the Preferred Path feature and to alternate between controllers for the preferred paths. We also concluded that striping should be across all HBAs.

Figure 5: Two disk groups with Preferred Path feature of Secure Path with no zoning

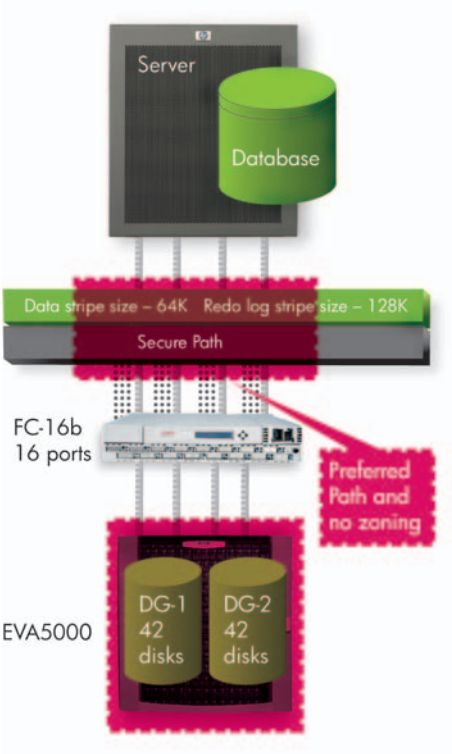

Table 5: Impact of LVM striping in a single disk group compared to LVM striping across multiple disk groups

	1 disk group		2 disk groups	
	Medium	Heavy	Medium	Heavy
Transactional throughput/m	20,426	27,162	20,654	27,316
Log file parallel write (ms)	0-1	1-2	0-1	1-2
Db file sequential read (ms)	11	15	11	15
I/O throughput/sec	6,249	8,513	6,413	8,554

Test series 2 – two disk groups

In a second set of tests, we evaluated the performance of HP StorageWorks EVA with two disk groups as compared with a single disk group. (See Figure 5.) The test configuration included Secure Path with the Preferred Path feature, no zoning and with VRAID 1 data files and redo logs.

LVM striping in a single disk group vs. multiple disk groups (two in our tests)

In the first test of this configuration, we evaluated double striping over the same spindles when LVM striping is used (single disk group). In the case of multiple disk groups, the Vdisks were selected in a round robin fashion from each disk group when creating a volume group, and any logical volumes were expended across all disk groups. The purpose of this test was to determine the impact of LVM striping in a single disk group compared to LVM striping across multiple disk groups. The results are presented in Table 5.

From these results, we determined that the number of disk groups does not affect performance when all of the data and logs are striped across all the disk groups—that there is no penalty for double striping.

Table 6: Performance of the HP StorageWorks EVA with VRAID 1 vs. VRAID 5

	Data/Logs—VRAID 1		Data/Logs—VRAID 5	
	Medium	Heavy	Medium	Heavy
Transactional throughput/m	20,654	27,316	18,399	24,123
Log file parallel write (ms)	0-1	1-2	0-1	1-2
Db file sequential read (ms)	11	15	12	18
I/O throughput/sec	6,413	8,554	5,703	7,604

Table 7: Performance of a single disk group vs. multiple disk groups

	1 disk group		2 disk groups	
	Medium	Heavy	Medium	Heavy
Transactional throughput/m	20,426	27,162	16,297	20,282
Log file parallel write (ms)	0-1	1-2	0-1	1-2
Db file sequential read (ms)	11	15	14	23
I/O throughput/sec	6,249	8,513	4,864	6,068

Comparing VRAID 1 and VRAID 5
In a variation of the multiple disk group tests, we compared performance of the HP StorageWorks EVA with VRAID 1 and VRAID 5. Data and logs were striped across all disk groups, and the same number of spindles was used in both test cases.

There are, however, performance implications for parity-based protection such as VRAID 5, especially for small writes. Data modification with small writes require four I/O operations: read old data, read old parity, write new parity and write new data. This is known as read-modify-write. During read-modify-write, read operations can be performed concurrently, then the write operations can be performed concurrently after the parity calculation. The results of this test are shown in Table 6.

As the data show, VRAID 1 outperformed RAID 5 significantly, showing a 12 percent degradation of transactional throughput for VRAID 5 compared to VRAID 1. It is important to note that we were able to achieve this result by allocating the same number of disks for both cases. The number of disks allocated and the array cache size for VRAID 5 configuration will determine the performance differences with a VRAID 1. Also, in VRAID 5, read-modify-write will affect both I/O and transactional throughput, so depending on write intensity and array de-staging, efficiency performance may vary. Ultimately, the trade-off between VRAID 1 and VRAID 5 is between capacity and performance.

Single disk group compared with multiple disk groups
In the final test, we directly compared performance of a single disk group with multiple disk groups, where the data files and redo logs are striped across different disk groups. This is usually referred to as traditional method by opposition to the Oracle SAME methodology. The results of this test are shown in Table 7.

The results of this test indicate that using multiple disk groups while Oracle files are striped across different disk groups can result in unbalanced I/O processing, which would have a serious impact on application performance. In this specific case, we observed degradation in transactional throughput of 25 percent under heavy load. While performance can change if application behavior changes, for most applications, there should be no reason to create multiple disk groups for performance purposes. For some applications in which the redo logs sync time is very critical, it may be beneficial to dedicate separate disk resources. Also in some cases, multiple disk groups may be viewed as a means to achieve high availability (to isolate logs from data), but since the EVA has no single point of failure, high availability is built in and there is no need for multiple disk groups.

Conclusion

Because performance is such a critical issue for Oracle DBAs, it is important to have a storage platform that can deliver consistently high performance, while providing the flexibility to support a range of application scenarios. As we have seen in this extensive set of tests, the HP StorageWorks EVA provides both performance and flexibility in supporting demanding transactional workloads.

In these tests, we could demonstrate that transactional throughput does not vary significantly between logs on VRAID 1, VRAID 5 and VRAID 0. We also determined that little variation in transactional throughput occurred whether read cache was enabled or disabled.

Because the overhead associated with HP StorageWorks Secure Path software was negligible, we recommend that the HP StorageWorks EVA be used with Secure Path software with the Preferred Path feature enabled.

Finally, if use of multiple disk groups is desired, it is best to stripe data files and logs across all the disk groups, rather than different disk groups, which can result in unbalanced I/O and degraded performance.

For most applications using Oracle 10*g*, the HP StorageWorks EVA configured with a single disk group, no zoning and using the Preferred Path feature of HP StorageWorks Secure Path software is the ideal solution to ensure maximum performance and flexibility.

For more information on how HP StorageWorks EVA supports Oracle 10*g*, please contact your HP representative or visit www.hp.com/go/oracle.

5983-0322EN, 1/2005