SAP® Sybase® IQ
on VMware vSphere®

November 2013

DEPLOYMENT AND TECHNICAL CONSIDERATIONS GUIDE

vmware®

Contents

Introduction

With the trend moving toward large logical CPU and high memory systems at cost-effective prices, more and more SAP® customers want to deploy their mission critical, high performance SAP Sybase® production databases on the VMware vSphere® virtual infrastructure. This includes SAP Sybase® IQ, a powerful, efficient column-based relational database management system (RDBMS) platform for extreme-scale enterprise data warehousing and big data analytics.

For this reason, VMware®, in cooperation with SAP, conducted testing to characterize the functionality of using SAP Sybase IQ 15.4, configured on VMware vSphere 5.1 virtual infrastructure. This paper describes the testing that was performed, the workload that was used, and the results that were observed. It also outlines the best practices for deploying this solution in the datacenter.

The paper is written for experienced architects and engineers who are responsible for SAP Sybase IQ and the VMware vSphere environment. This paper assumes that the reader has advanced knowledge of SAP Sybase products and the VMware vSphere virtualization and cloud computing platform.

Analytic Workloads

Analytic workloads are the perfect type of workloads for virtualization. The processing of analytic workloads is characterized by its peaks and valleys. The peaks are fairly predictable events that occur at end of the month, end of the quarter, or at end of year. Peak periods are followed by quiet times where resource consumption and utilization are well below the capabilities of the physical servers on which these workloads are deployed. During these valleys, these resources are locked into the physical servers and wasted until the next predictable peak processing period is encountered.

When deploying analytic environments, architects must size and accommodate for peak activity. This typically entails throwing "big iron" at the solution. This requires large capital expenditures for systems that ultimately results in a tremendous amount of idle resources.

Value of Running Analytic Workloads on the vCloud Suite

These are just some of the reasons that businesses choose to run analytic workloads on the VMware® vCloud Suite®. With VMware technologies, resources can be reserved to guarantee performance at peak processing periods in order to maintain enterprise service level agreements (SLAs). After peak processing has completed, resources can then be released and allocated elsewhere in the infrastructure.

Freeing Sybase IQ Compute Resources From Physical Servers

In many cases, Sybase IQ can run and process queries hundreds or thousands of times faster than traditional relational databases. When Sybase IQ runs on physical hardware, it might actually heighten the unwanted side effect of idle resources associated with analytic workloads. This is because queries are being completed at a significantly faster rate.

However, running Sybase IQ on vSphere enables SAP application owners and VI-Administrators to redeploy idle resources elsewhere in their enterprise, essentially freeing Sybase IQ from a physical instantiation. Using Sybase IO on vSphere can significantly reduce the overall hardware capital expenditure because these resources are not locked into any one physical server. They can be amortized over many instances or applications.

Solution Overview

VMware and SAP Sybase jointly conducted testing to characterize the functionality of using SAP Sybase IQ 15.4, configured on VMware vSphere 5.1 virtual infrastructure. The SAP Sybase and VMware vSphere products used in this solution are described in the sections below.
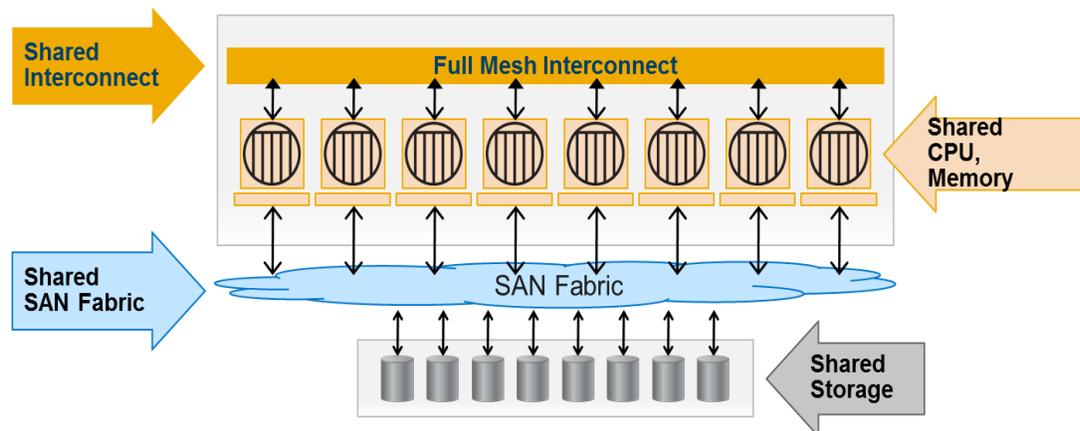
SAP Sybase IQ

SAP Sybase IQ is a powerful, efficient column-based relational database management system (RDBMS) platform for extreme-scale enterprise data warehousing and big data analytics. This includes Hadoop integration/federation, in-database advanced analytics, and the native MapReduce programming framework.

SAP Sybase IQ delivers major innovations in the areas of performance optimization, data compression, and data protection. The SAP Sybase IQ platform is used by businesses that want to exploit the value of massive amounts of data for more accurate insight into business performance and market dynamics. This data allows businesses to act on insights that were previously impossible or impractical to analyze.

The SAP Sybase IQ columnar architecture enables scale-out to multiple nodes in a shared disk environment. This enables easy, cost-effective enterprise data warehouse expansion with minimal hardware and database administration requirements, and without the overhead costs of the traditional Massively Parallel Processor (MPP) architecture.

Figure 1. SAP Sybase IQ Multiplex architecture

The SAP Sybase IQ Multiplex grid architecture is an architectural framework that enables an elastic, massively parallel, high performance enterprise data warehousing infrastructure that is based on a combination of a powerful Shared Everything MPP and a rich application services layer.

The SAP Sybase IQ Multiplex Grid option enables the Shared Everything MPP layer that delivers the full power of massively parallel processing by enabling incremental scale-out of concurrent query jobs, load jobs, or both in an elastic and cost effective manner for high performance enterprise data warehouses.

Figure 2. SAP Sybase IQ Multiplex grid architecture
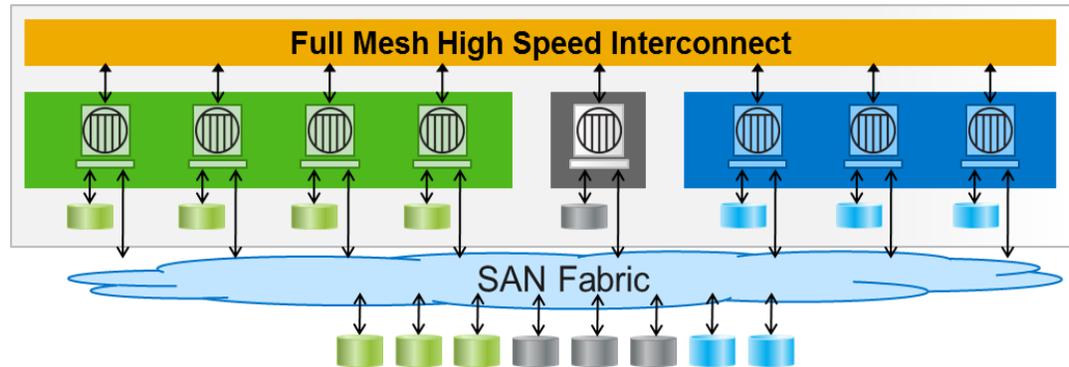
For more information, see the "Resources" section at the end of this paper.

VMware vSphere

VMware vSphere is an optimal virtualization platform and an enabler of cloud computing architectures. vSphere allows IT to meet service level agreements (SLAs) for the most demanding business-critical applications at the lowest total cost of ownership (TCO). VMware vSphere delivers control over all IT resources with the highest efficiency and choice in the industry.

VMware vSphere virtualization solutions provide for:

- **Consolidation**. VMware virtualization allows multiple application servers to be consolidated onto one physical server, with little or no decrease in overall performance. This helps to minimize or eliminate underutilized server hardware, software, and infrastructure.

- **Manageability**. The live migration of virtual machines from server to server and its associated storage is performed with no downtime using VMware vSphere® vMotion®, which simplifies common operations such as hardware maintenance, and VMware vSphere® Storage vMotion®.

- **Availability**. High availability can be enabled to reduce unplanned downtime and enable higher service levels for applications. VMware vSphere® High Availability (HA) ensures that, in the event of an unplanned hardware failure, the affected virtual machines are automatically restarted on another host in a VMware cluster.

- **Automation**. VMware automated load balancing takes advantage of vMotion and Storage vMotion to migrate virtual machines among a set of VMware® ESXi® hosts. VMware vSphere® Distributed Resource Scheduler™ (DRS) and VMware vSphere® Storage DRS™ enable automatic resource relocation and optimization decisions for virtual machines and storage.

- **Provisioning**. VMware virtualization encapsulates an application into an image that can be duplicated or moved, greatly reducing the cost of application provisioning and deployment.

Figure 3. VMware vSphere virtual infrastructure
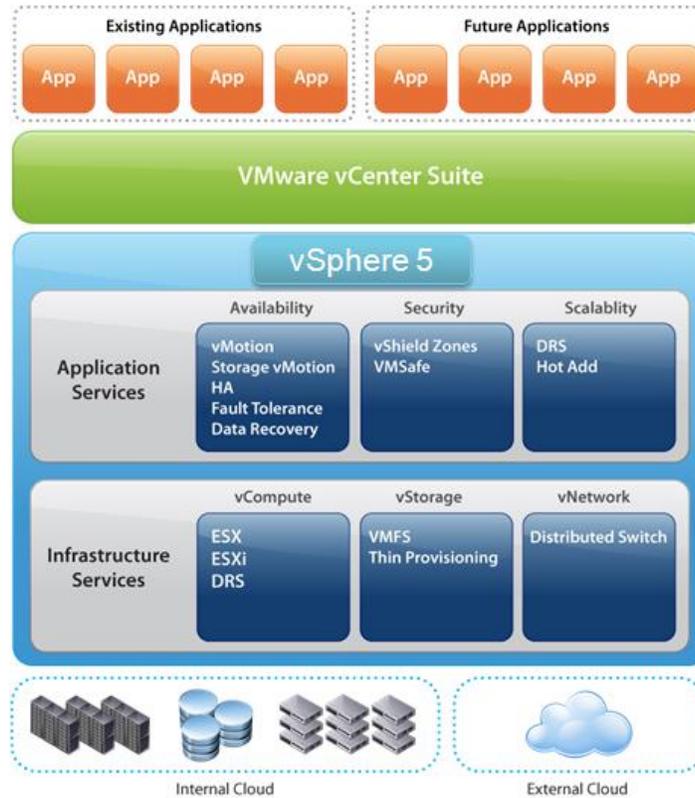
vSphere creates a layer of abstraction between the resources required by an application, the operating system, and the underlying hardware that provides those resources. vSphere enables multiple, isolated execution environments to share a single hardware platform. It implements each environment with its own set of hardware resources.

For more information, see the "Resources" section at the end of this paper.

vMotion

VMware vMotion enables the live migration of running virtual machines from one physical server to another with zero downtime, continuous service availability, and complete transaction integrity. vMotion is a key enabling technology for creating the dynamic, automated, and self-optimizing datacenter. This capability makes hardware maintenance possible at any time, and vMotion does not require clustering or redundant servers, as described below.

Figure 4. VMware vMotion
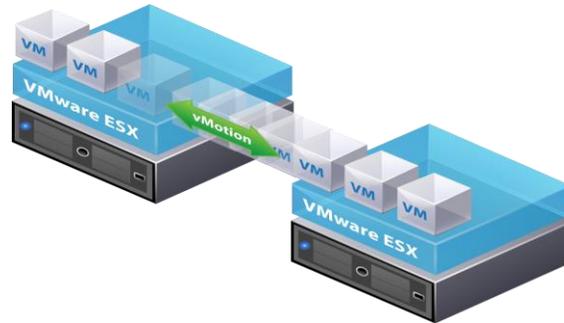
VMware vMotion:

- Moves entire running virtual machines instantly. Performs live migrations with zero downtime, undetectable to the user.

- Manages and schedules live migrations with ease at pre-defined times without an administrator's presence, with the reliability and manageability that is derived from a production-proven product.

- Performs multiple concurrent migrations of a virtual machine running any operating system, across any type of hardware and storage that is supported by vSphere, complete with an audit trail.

- Moves online workloads from one ESXi Server host machine to another in order to maintain service levels and performance goals.

- Continuously and automatically optimizes virtual machine placement within resource pools. Proactively moves virtual machines away from failing or underperforming servers.

- Performs hardware maintenance without the need to schedule downtime and disrupt business operations.

The entire state of a virtual machine is encapsulated by a set of files stored on shared storage, such as Fibre Channel or iSCSI Storage Area Network (SAN) or Network Attached Storage (NAS). VMware storage Virtual Machine File System (VMFS) allows multiple installations of VMware ESXi® to access the same virtual machine files concurrently.

The active memory and precise execution state of the virtual machine is rapidly transferred over a high speed network, allowing the virtual machine to instantaneously switch from running on the source ESXi host to the destination ESXi host. vMotion keeps the transfer period imperceptible to users by keeping track of on-going memory transactions in a bitmap.

Once the entire memory and system state has been copied over to the target ESXi host, vMotion suspends the source virtual machine, copies the bitmap to the target ESXi host, and resumes the virtual machine on the target ESXi host. The networks being used by the virtual machine are also virtualized by the underlying ESXi host, ensuring that even after the migration, the virtual machine network identity and network connections are preserved. vMotion manages the virtual MAC address as part of the process. Once the destination machine is activated, vMotion pings the network router to ensure that it is aware of the new physical location of the virtual MAC address. Since the migration of a virtual machine with vMotion preserves the precise execution state, the network identity, and the active network connections, the result is zero downtime and no disruption to users.

Testing Architecture

The SAP Sybase IQ testing with VMware virtual infrastructure was conducted in the VMware Lab in Palo Alto, California using the testing configuration shown below.

Figure 5. SAP Sybase IQ testing configuration

20 x 35 GB and 1 x 400 GB drives configured on
EMC VNX 5700 FC Storage
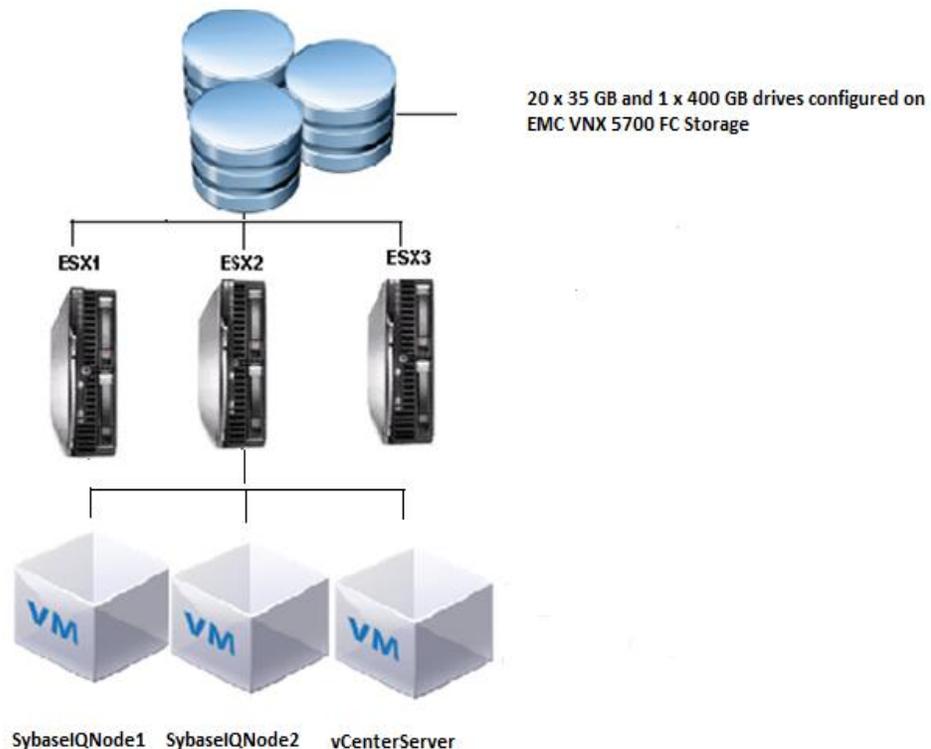
The architecture for this testing included:

- Hewlett-Packard® servers included 3 HP® Proliant® DL 380p Gen8 Servers, with one ESXi host on each physical server (Host 1, Host 2, and Host 3)

- EMC® VNX® 5700 FC Unified Storage System

- One virtual machine was configured on each host as each Sybase IQ node was configured with 16 vCPUs, and the CPU utilization on the VM was expected to be high. This included one virtual machine for the SAP Sybase IQ database: Sybase IQ Node 1 (Host 1), Sybase IQ Node 2 (Host 2), and the VMware® vCenter Server™ appliance (Host 3).

For more details about the set up of the ESXi host servers and virtual machines used to perform the tests, see the "Appendix A: Testing Configuration" section later in this paper.

Testing

Because SAP Sybase supports the VMware vSphere platform, many customers are very interested in deploying their solution on VMware. However, given the mission-critical nature of the application suite, customers need assurance that the combined solution provides the required functionality. For this reason, VMware and SAP Sybase jointly conducted testing in order to characterize the functionality of using SAP Sybase IQ 15.4, configured on VMware vSphere 5.1 virtual infrastructure.

The testing methodology, test scenarios, test set, workload used, results observed, and best practices recommendations are described in the sections below.

Testing Methodology

The testing methodology was to characterize the functionality and performance of running SAP Sybase IQ 15.4 on VMware vSphere 5.1 virtual infrastructure for a given workload. The testing was also used to evaluate the operation of VMware vMotion with SAP Sybase deployments that run in a virtual environment.

Virtual Machine File System

VMware vSphere VMFS provides high-performance, clustered storage virtualization that is optimized for virtual machines. With VMFS, each virtual machine is encapsulated into a small set of files. VMFS is the default storage management interface used to access those files on physical SCSI disks and partitions. VMFS allows multiple ESXi instances to access shared virtual machine storage concurrently. It also enables virtualization based, distributed infrastructure services such as vMotion, VMware DRS, and VMware High Availability to operate across a cluster of ESXi hosts.

In order to balance performance and manageability in a virtual environment, it is an accepted best practice to deploy databases using VMFS. Raw Device Mappings (RDMs) are sometimes erroneously selected to provide increased performance. In contrast, the two dominant workloads associated with databases, random read/write and sequential writes, have nearly identical performance throughput characteristics when deployed on VMFS or using RDM.

The testing done using RDMs provides conclusive evidence database workloads should be run on VMFS.

Test Scenarios

The test scenarios used the Single-node and Multiplex configurations as shown in the table below. The results for each of these tests are described in the sections below.

Table 1. Types of test scenarios

Test scenario	Tests case
Single-node configuration	• **Duration**: Time required to load the TPC-H data. • **Power test**: Execute a single run-through of all of the tests in a stream. • **Throughput test**: Run all of the streams in parallel. Once started, all streams follow their query sequence independently, without waiting for the queries from another stream to finish.
Multiplex configuration	• **Power test**: Test the second, reader-node of the Multiplex to allow distributed query processing (DQP) under the grid architecture. • **Throughput test**: Test the second, reader-node of the Multiplex to allow DQP under the grid architecture.
Both configurations	• **vMotion test**: Run the power and throughput tests for vMotion on IQNode1 and IQNode2 with both the Single-node and Multiplex configurations.

Test Set

The test set for each testing scenario is described in the sections below.

Virtual Tests

Table 2. Virtual tests: Test scenarios and test cases

Test scenario	Test case
Single-node and Multiplex (with VMFS and Multiwriter enabled)	• Load test • Power test • Throughput test
Single-node and Multiplex (with RDM and SCSI Sharing)	• Load test • Power test • Throughput test

Physical Tests

Table 3. Physical tests: Test scenario and test cases

Test scenario	Test case
Single-node and Multiplex	• Load test • Power test • Throughput test

Workload Used

An unaudited 300 GB TPC-H environment was configured and used as a standardized test platform. The same 300 GB of generated input data was used for all load tests. The virtual machines had built-in load drivers. The same six query streams were used for the power and throughput tests. The entire TPC-H benchmark test was run to measure performance against a comparable non-virtualized environment.

Table 4. Workload scripts and execution

Script name	Frequency of execution
run_tpch load	Load all the data for the TPC-H tables.
run_tpch power	Perform a single run across all tests in a stream (with no refresh).
run_tpch throughput	Run all streams in parallel (with no refresh).

Results Observed

The results observed for each test case are described in the sections below. Overall, the testing results were very positive. VMware functioned well and all of the tests passed successfully. The results of the tests indicate that SAP Sybase IQ can be configured to run in a Single-node or Multiplex configuration under VMware both with VMDK and Multiwriter enabled, and with RDM and SCSI Sharing.

Testing was performed to provide detailed information for resource utilization including CPU, memory, and IOPS in order to obtain the proper sizing and capacity planning. The metrics for each of these tests is provided in the graphs below.

The test results that were observed and their impact are summarized in the table below. For more detailed information, review each specified test area in the sections below.

Table 5. Testing results summary

Test case	Test results
Load	All of the lines items load successfully without any failure.
Power	The test ran successfully with the single run running through all tests without any problems. The query execution time was at par with the physical tests.
Throughput	The parallel processing stream test worked as expected with the query processing times close to physical benchmarks.
vMotion	vMotion tested successfully with both the Single-node and Multiplex configurations. Both IQNodes were tested with vMotion without issue.

Virtual Tests

Single-node with VMDK and Multiwriter Enabled

Load Test

Table 6. Application level metrics

Type of metric	Definition
Database load time	38mins:30secs (2310 secs, 9.11 MB/s/core)
Database load size	328 GB
Number of vCPUs	16

Table 7. Timing breakup

Line item	Region	Nation	Customer	Orders	Supplier	Part	Part supplier
1514	0	1	57	452	6	43	237

Power Test

Table 8. Stream execution duration

StreamID	Start date	Start time	End date	End time	Duration
stream0	2013-03-08	14:24:31	2013-03-08	14:38:32	14mins:1sec

Throughput Test

Table 9. Numerical quantities

Type of metric	Definition
Database scale factor	300 GB
Query streams for throughput test	6
TPC-H throughput@300 GB	43891.6
TPC-H throughput interval	3248 seconds

Figure 6. Single-node: Load, power, and throughput tests (CPU Utilization for VMDK with Multiwriter)

VMDK Multiwriter Single Node CPU Utilization

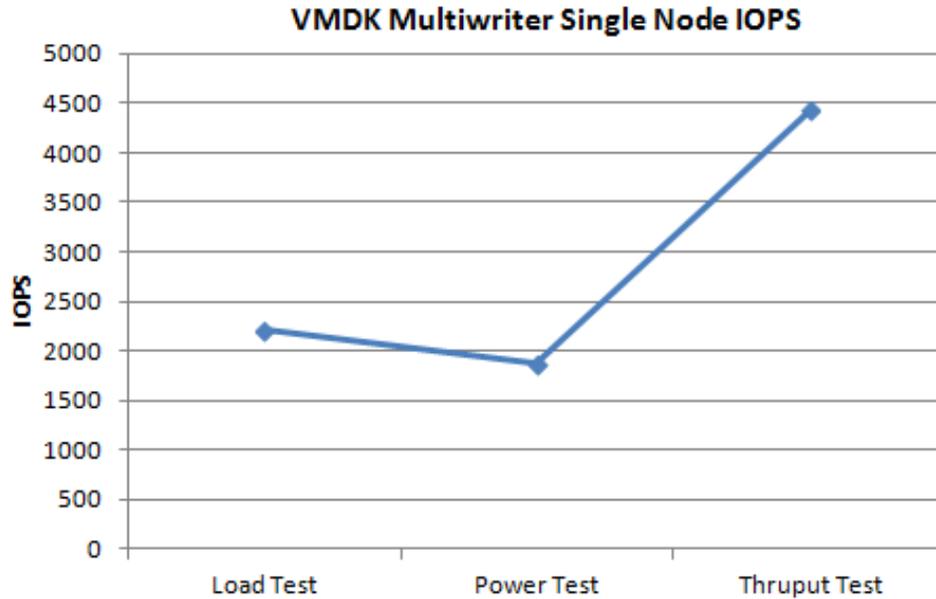
Figure 7. Single-node: Load, power, and throughput tests (memory utilization for with VMDK Multiwriter)

VMDK Multiwriter Single Node Memory Utilization

Figure 8. Single-node: Load, power, and throughput tests (IOPS for VMDK with Multiwriter)

Multiplex with VMDK and Multiwriter Enabled

Load Test

Table 10. Application level metrics

Type of metric	Definition
Database load time	38mins:37secs (2317 secs, 9.08 MB/s /core)
Database load size	328 GB
Number of CPUs	16 vCPUs

Table 11. Timing breakup

Line item	Region	Nation	Customer	Orders	Supplier	Part	Part supplier
1542	0	0	53	431	6	42	243

Power Test

Table 12. Stream execution duration

StreamID	Start date	Start time	End date	End time	Duration
stream0	2013-03-12	16:03:27	2013-03-12	16:16:33	13mins:6secs

Throughput Test

Table 13. Numerical quantities

Type of metric	Definition
Database scale factor	300 GB
Query streams for throughput test	6
TPC-H throughput@300 GB	44149
TPC-H throughput interval	3229 seconds

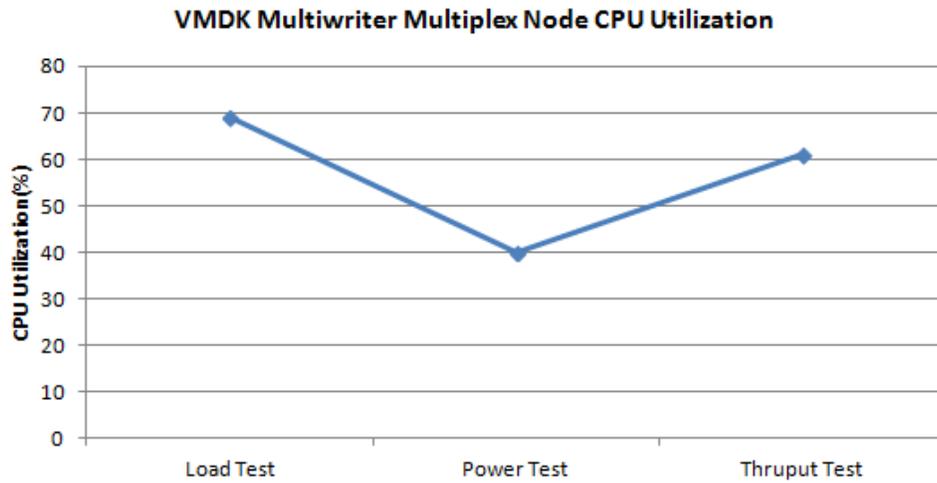**Figure 9. Multiplex: Load, power, and throughput tests (CPU utilization for VMDK with Multiwriter)**

VMDK Multiwriter Multiplex Node CPU Utilization

Figure 10. Multiplex: Load, power, and throughput tests (memory utilization for VMDK with Multiwriter)

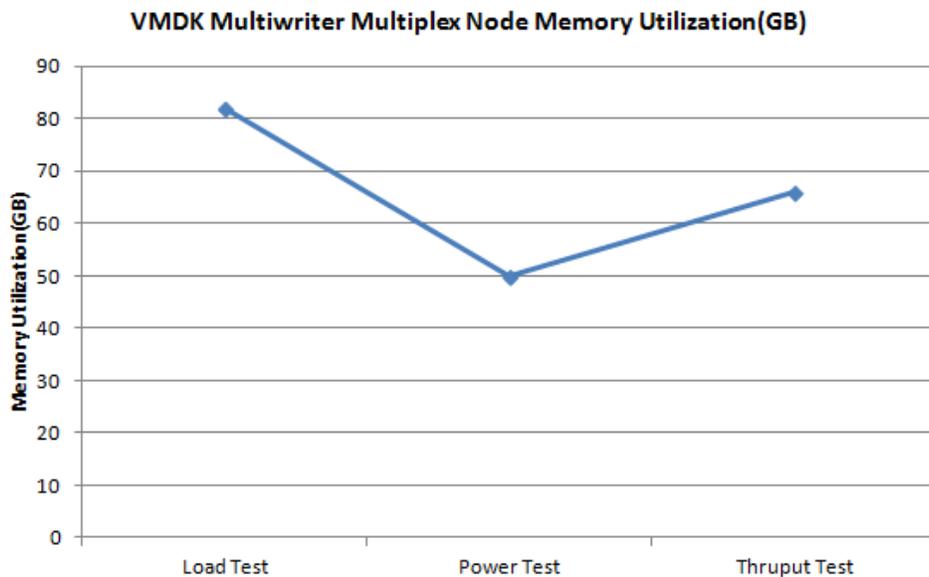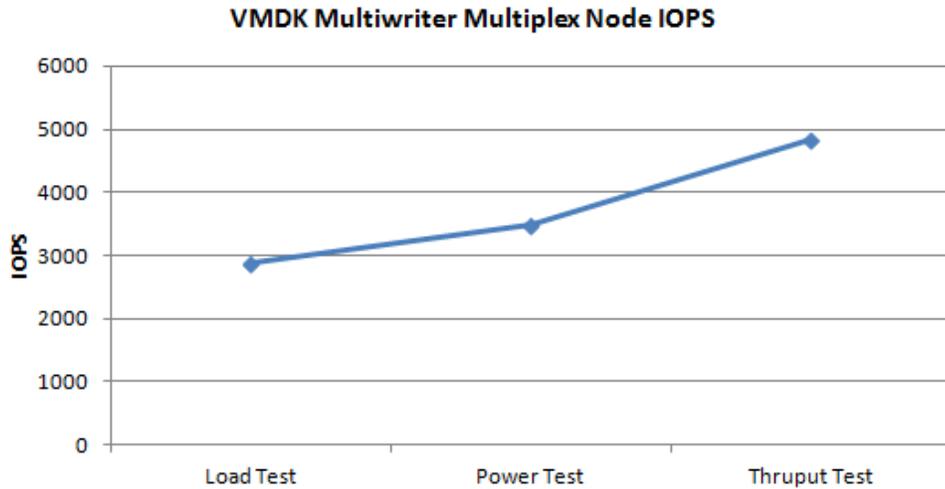**VMDK Multiwriter Multiplex Node Memory Utilization(GB)**

Figure 11. Multiplex: Load, power, and throughput tests (IOPS for VMDK with Multiwriter)

Single-node with RDM and SCSI Sharing

Load Test

Table 14. Application level metrics

Type of metric	Definition
Start database load time	2013-03-01 17:07:20
End database load time	2013-03-01 17:46:34
Database load time	39mins:14secs (2354 secs, 9.86 MB/s/core)
Database load size	328 GB

Power Test

Table 15. Stream execution duration

StreamID	Start date	Start time	End date	End time	Duration
stream0	2013-03-01	10:47:27	2013-03-01	11:00:34	13mins:7secs

Throughput Test

Table 16. Numerical quantities

Type of metric	Definition
Database scale factor	300 GB
Query streams for throughput test	6
TPC-H throughput@300 GB	43516.2
TPC-H throughput interval	3276 seconds

Figure 12. Single-node: Load, power, and throughput tests (CPU utilization for RDM SCSI with multi-sharing)
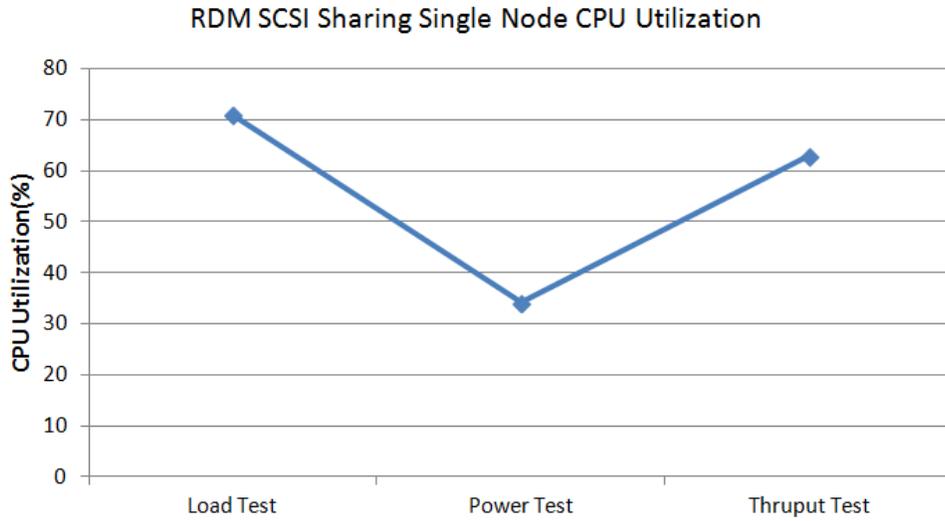

Figure 13. Single-node: Load, power, and throughput tests (memory utilization for RDM SCSI with multi-sharing)
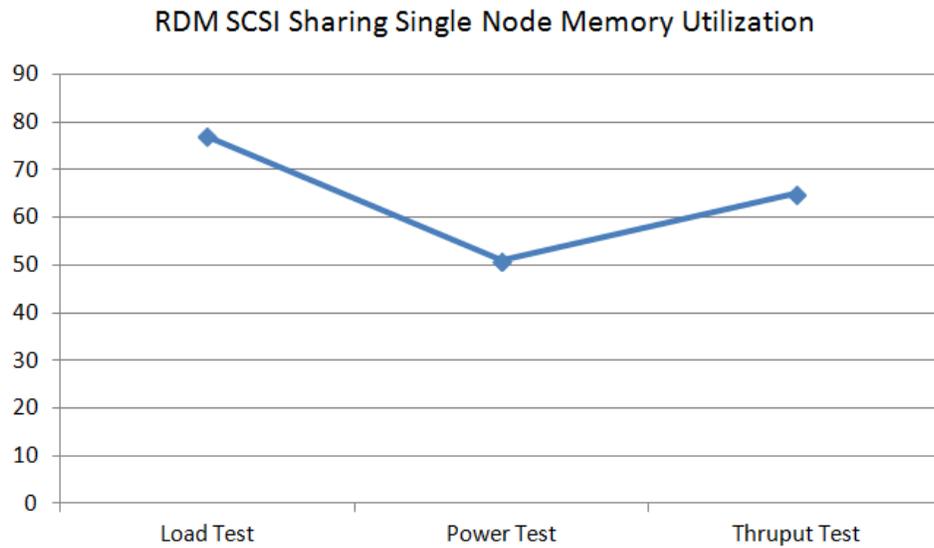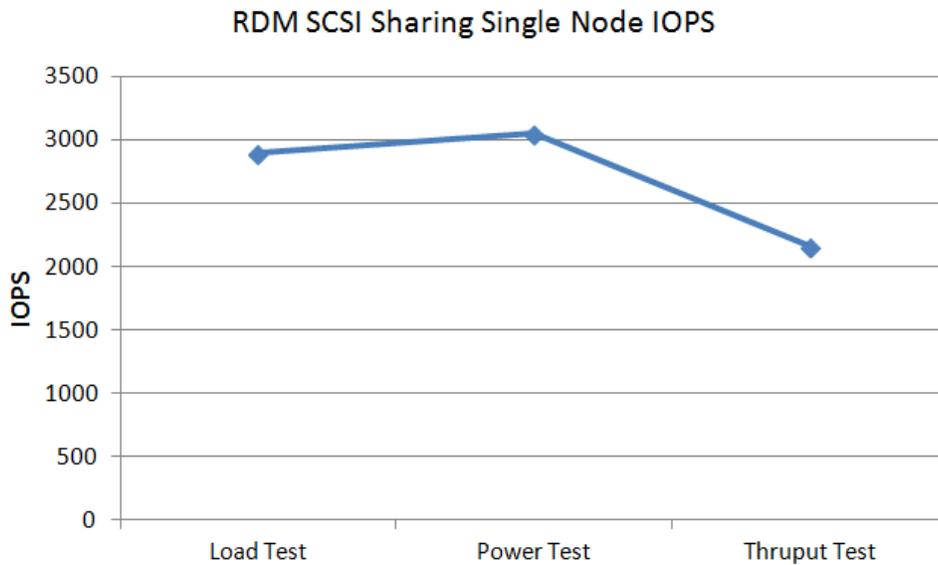
Figure 14. Single-node: Load, power, and throughput tests (IOPS for RDM SCSI with multi-sharing)

Multiplex with RDM and SCSI Sharing

Load Test

Table 17. Application level metrics

Type of metric	Definition
Start database load time	2013-03-01 15:29:11
End database load time	2013-03-01 16:08:36
Database load time	39mins:25secs (2365 secs, 8.90 MB/s/core)

Power Test

Table 18. Stream execution duration

StreamID	Start date	Start time	End date	End time	Duration
stream0	2013-03-04	13:28:37	2013-03-04	13:42:42	14mins:5secs

Throughput Test

Table 19. Numerical quantities

Type of metric	Definition
Database scale factor	300 GB
Query streams for throughput test	6
TPC-H throughput@300 GB	43891.7
TPC-H throughput interval	3248 seconds

Figure 15. Multiplex: Load, power, and throughput tests (CPU utilization for RDM SCSI with multi-sharing)
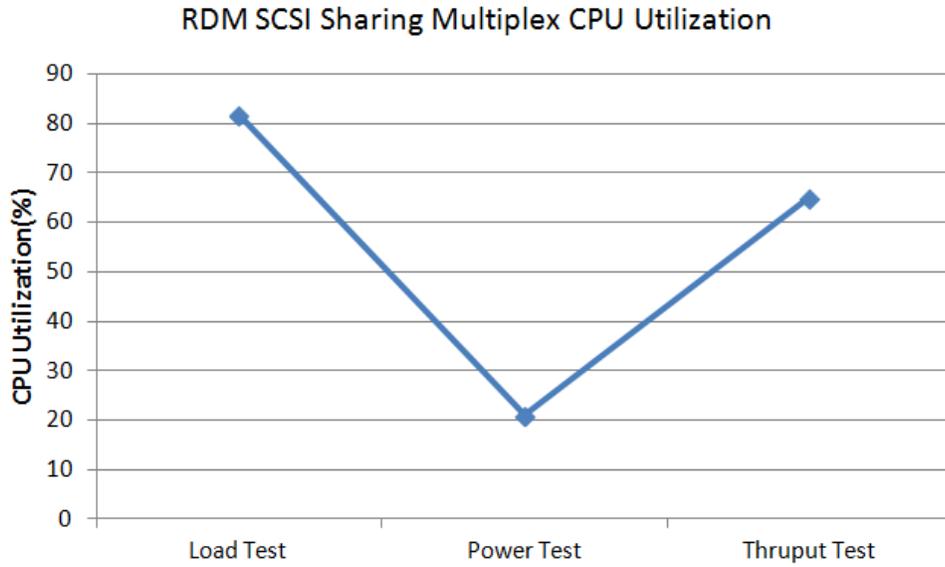

Figure 16. Multiplex: Load, power, and throughput tests (memory utilization for RDM SCSI with multi-sharing)
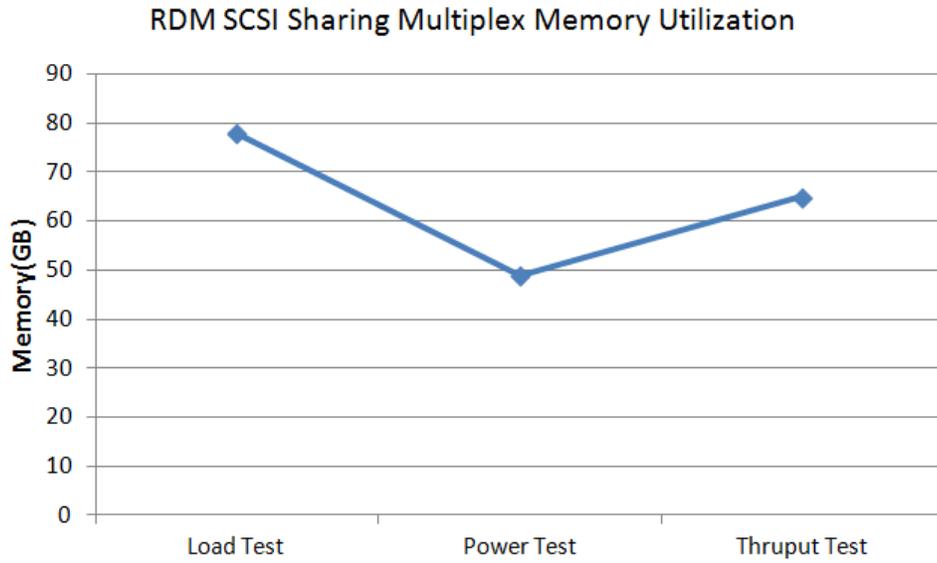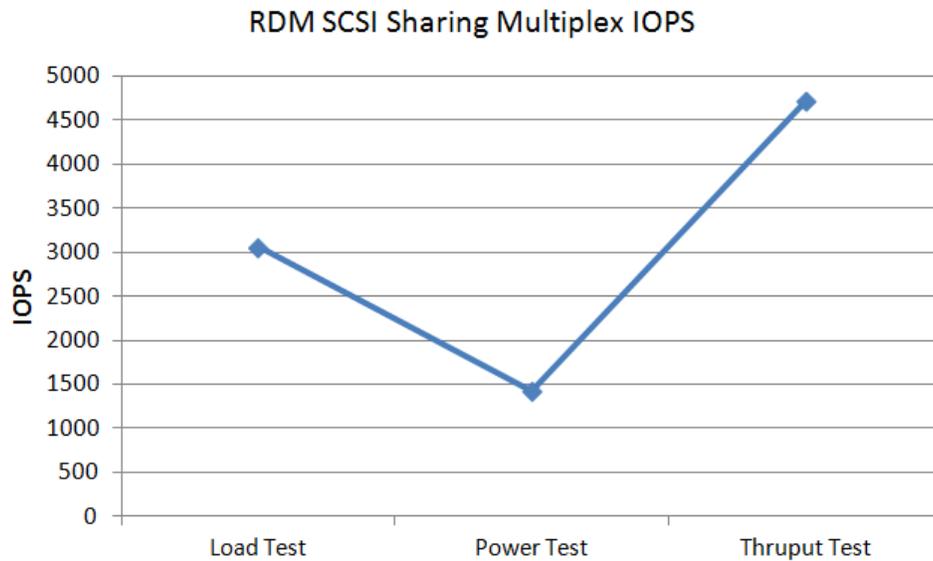
Figure 17. Multiplex: Load, power, and throughput tests (IOPS for RDM SCSI with multi-sharing)

Physical Tests

Single-node

Load Test

Table 20. Application level metrics

Type of metric	Definition
Start database load time	2013-02-27 04:31:40
End database load time	2013-02-27 05:06:17
Database load time	34mins:37secs (2077 secs, 10.13 MB/s/core)
Database load size	328 GB
Number of CPUs available	16

Power Test

Table 21. Stream execution duration

StreamID	Start date	Start time	End date	End time	Duration
stream0	2013-02-27	13:42:30	2013-02-27	13:55:32	13mins:2secs

Throughput Test

Table 22. Application level metrics

Type of metric	Definition
Database scale factor	300 GB
Query streams for throughput test	6

Type of metric	Definition
TPC-H throughput@300 GB	45913.0
TPC-H throughput interval	3105 seconds

Multiplex

Power Test

Table 23. Stream execution duration

StreamID	Start date	Start time	End date	End time	Duration
stream0	22013-02-28	09:31:38	2013-02-28	09:44:40	13mins:2secs

Throughput Test

Table 24. Application level metrics

Type of metric	Definition
Database scale factor	300 GB
Query streams for throughput test	6
TPC-H throughput@300 GB	44149.9
TPC-H throughput interval	3229 seconds

Timing Comparison

The timing comparison between the different types of runs is shown in the figure below.

Figure 18. Timing comparison
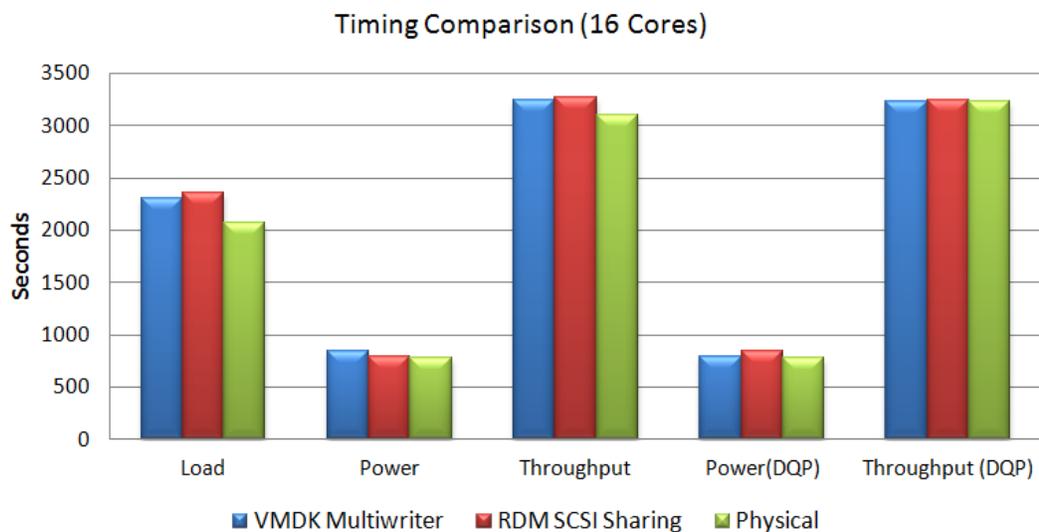
The timing comparison shows little to no difference in the timing between the use of VMDK in Multiwriter mode and with RDM SCSI Sharing.

For loading, 10 to 15 percent overhead from the physical timings was observed, while an 0 to 8 percent overhead from the physical environment was observed for the power and throughput tests using the Single-node or Dual-node DQP configuration.

The difference between loading and query performance is due to the ability of SAP Sybase IQ to cache pages in the node memory for better query performance. In this testing, loading was performed against an empty database each time, reducing the effectiveness of caching.

The same physical disk configuration was used for all of the testing, including virtual VMDK with Multiwriter enabled, virtual RDM with SCSI sharing, and physical testing. As mentioned previously, a set of functional tests were performed only. This testing did not involve a "reference architecture" storage configuration, according to SAP Sybase IQ best practices, to improve performance.

Given the physical disk configuration, the virtual servers were able to provide performance on par with the physical servers.

vMotion Testing

Single-node vMotion

The results for the power and throughput tests demonstrate that the vMotion operation was completed in just under three minutes, as shown below.

Power test

Table 25. Stream execution duration

StreamID	Start date	Start time	End date	End time	Duration
stream0	2013-03-15	17:02:25	2013-03-15	17:16:34	14mins:9secs

Throughput Test

Table 26. Numerical quantities

Type of Metric	Definition
Database scale factor	300 GB
Query streams for throughput test	6
TPC-H throughput@300 GB	44957.4
TPC-H throughput interval	3229 seconds

Multiplex vMotion

The results for the power and throughput tests demonstrate that the vMotion operation was completed in just under three minutes, as shown below.

Power Test

Table 27. Stream execution duration

StreamID	Start date	Start time	End date	End time	Duration
stream0	2013-03-18	15:30:55	2013-03-18	15:45:02	14mins:7secs

Throughout Test

Table 28. Application level metrics

Type of metric	Definition
Database scale factor	300 GB
Query streams for throughput test	6
TPC-H throughput@300 GB	44787.9
TPC-H throughput interval	3183 seconds

Deployment Best Practices

Based on this testing, the recommended best practices for this solution include:

1. When using Intel-based systems, ensure that the BIOS settings enable the Hyperthreading, VT, and EPT options on all ESXi hosts.

2. Use Eager zeroed disks for better I/O performance by using the VMware vmkfstools command:

   vmkfstools –w abc.vmdk

3. Decrease disk latency by using fibre-channel adapters and configure the proper (redundant array of independent disks) RAID configuration with sufficient spindles.

4. Install VMware Tools on the virtual machines. The VMware Tools package provides support for shared folders and for drag-and-drop operations. Other tools in the package support time synchronization of the guest operating system with time on the host, automatic grabbing and releasing of the mouse cursor, copying and pasting between guest and host, and improved mouse performance in some guest operating systems.

5. Consider using server-class network interface cards (NICs) for the best performance and configure paravirtualized vmxnet3 adapters for better network throughput.

6. Configure Sybase IQ using VMDKs and Multi-writer flag. The use of RDMs limit VMware features and functionality.

7. Disconnect or disable unused or unnecessary physical hardware devices to free interrupt resources. These hardware devices include:

    ●   COM ports

    ●   LPT ports

    ●   USB controllers

    ●   Floppy drives

    ●   Optical drives (CD or DVD drives)

    Traditionally, some devices such as USB controllers operate on a polling scheme that consumes extra CPU resources. In addition, some PCI devices reserve blocks of memory, making that memory unavailable to ESXi.

Conclusion

VMware, in cooperation with SAP, conducted testing to characterize the functionality of using SAP Sybase IQ 15.4, configured on VMware vSphere 5.1 virtual infrastructure. This paper describes the testing that was performed, the workload that was used, and the results that were observed. It also outlines the best practices for deploying this solution in the datacenter.

Overall, the testing results were very positive. The VMware virtual infrastructure functioned well and all of the tests passed successfully. The results of the tests indicate that SAP Sybase IQ should be configured to run in a Single-node or Multiplex configuration with VMDKs and Multiwriter enabled. The paper also provides conclusive evidence showing that VMDKs and RDMs have nearly identical performance characteristics for the two dominant workloads associated with databases.

These results indicate that running SAP Sybase IQ in a virtual environment using VMware vSphere virtual infrastructure has the potential to reduce cost, increase service levels, and simplify the manageability of the application. Additional testing is planned to configure disk and new server virtualization capabilities in order to demonstrate the performance of SAP Sybase IQ.

Appendix A: Testing Configuration

Hardware and Storage Configuration

Tables A-1 describes the configuration of VMware ESXi host servers and storage used in this testing.

Table A-1. ESXi host hardware configuration

Hardware	Configuration
Server	• 3 HP Proliant DL 380p Gen8 Servers. Each server is equipped with: • Intel(R) Xeon(R) CPU E5-2680 0 @ 2.70 GHz • 2 sockets • 8 cores per socket • 256 GB RAM • 2 x 1 GB NICS per ESXi – Broadcom 5706 • 2 x 10 GB NICS per ESXi – Broadcom 57810 • Hyperthreading: Enabled
Storage	• EMC VNX 5700 FC Unified Storage System equipped with: • 9 Solid State Drives (SSDs) • 20 x 35 GB partitions for IQ devices • 1 x 400 GB partition for source data file system

Installed Software

Table A-2 lists the installed software used in this testing.

Table A-2. Software Installed

Software Type	Software configuration
VMware	• VMware vSphere ESXi 5.1(Build 799733), with vMotion enabled • VMware vCenter Server 5.1
SAP Sybase	• Sybase IQ 15.4.0.3019 ESD 2

Virtual Machine Configuration

Table A-3 describes the configuration of virtual machines running on the ESXi host servers in this testing.

Table A-3. Virtual machine configurations

Virtual machine	Hardware configuration
Sybase IQ Node 1	16 vCPUs 120 GB memory - Ethernet card (10 Gbps network) (1 x 35 GB), (20 x 30 GB), and (1 X 395 GB) storage
Sybase IQ Node 2	16 vCPU 40 GB memory - Ethernet card (20 x 30 GB drives shared between both the nodes)
vCenter Server	2 vCPU 4 GB memory 1 Ethernet card (1 X 40 GB) storage

Linux Configuration

The Linux operating system was configured using the SAP Sybase IQ installation documentation. For more information, see the "Resources" section at the end of this paper.

SAP Sybase IQ Multiplex servers require shared read/write access at the raw device level.

The devices were partitions that served as a SAP Sybase IQ consistency group, and mapped to an EMC VNX 5700 with 9 SSDs configured to provide 20 x 35 GB devices. The devices were mapped to dbspace files as follows:

- 2 devices for shared IQ_SYSTEM_MAIN

- 12 devices for shared IQ_USER_MAIN

- 2 devices for shared IQ_SHARED_TEMP

- 2 devices for node 1 IQ_SYSTEM_TEMP

- 2 devices for node 2 IQ_SYSTEM_TEMP

- 1 device (450 GB) for ext4 filesystem used to host source load data (node 1)

The devices were configured to provide reasonable and consistent performance. They were not configured for optimal I/O performance (per SAP Sybase IQ sizing guidelines). The same devices were mounted to the SAP Sybase IQ servers for each phase of testing including physical, virtual with RDM sharing, and virtual with VMware VMDK sharing.

The block devices that were presented physically using RDM or VMDK Multiwriter were mapped to raw devices in the SAP Sybase IQ node operating system. Ownership was changed using major/minor block device numbers under /etc/udev/rules.d/60-raw.rules and /etc/udev/rules.d/99-raw.rules. For more information, see "Appendix B: Code Samples".

IQ Configuration/Multiplex Configuration

The parameters below were used to start both the coordinator and reader nodes for the test.

```
# tpch.cfg
# ------------------------------------------------------------
# Default startup parameters for the ASIQ demo database
# ------------------------------------------------------------
-n  tpch_15_4
-x  tcpip{port=5788}

-c 500M
-gd all
-gm 50

-gc 5000
-gr 5000
-gp 4096
-tl 0

-iqmc 40000
-iqtc 40000

-iqpartition 4
-iqgovern 20
-iqnumbercpus 14
```

Appendix B: Code Samples

In addition to the code samples, go to:

- Disabling simultaneous write protection provided by VMFS using the Multiwriter flag (1034165) (to enable Multiwriter):

  http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1034165

- Change the SCSI Bus Sharing Configuration (to enable SCSI Bus Sharing):

  http://pubs.vmware.com/vsphere-4-esx-vcenter/index.jsp?topic=/com.vmware.vsphere.vmadmin.doc_41/vsp_vm_guide/configuring_virtual_machines/t_change_the_scsi_controller_or_scsi_bus_sharing_configuration.html

For more information, see the "Resources" section at the end of this paper.

/etc/udev/rules.d/60-raw.rules

```
#/etc/udev/rules.d/60-raw.rules
# sybase-virutalnode1 / sybase-virtualnode2 raw device bindings
# IQ_SYSTEM_MAIN shared rw dbspace file devices
ACTION=="add", ENV{MAJOR}=="8", ENV{MINOR}=="16", RUN+="/bin/raw /dev/raw/raw1 %M %m"
ACTION=="add", ENV{MAJOR}=="8", ENV{MINOR}=="32", RUN+="/bin/raw /dev/raw/raw2 %M %m"
# IQ_USER_MAIN shared rw dbspace file devices
ACTION=="add", ENV{MAJOR}=="8", ENV{MINOR}=="48", RUN+="/bin/raw /dev/raw/raw3 %M %m"
ACTION=="add", ENV{MAJOR}=="8", ENV{MINOR}=="64", RUN+="/bin/raw /dev/raw/raw4 %M %m"
ACTION=="add", ENV{MAJOR}=="8", ENV{MINOR}=="80", RUN+="/bin/raw /dev/raw/raw5 %M %m"
ACTION=="add", ENV{MAJOR}=="8", ENV{MINOR}=="96", RUN+="/bin/raw /dev/raw/raw6 %M %m"
ACTION=="add", ENV{MAJOR}=="8", ENV{MINOR}=="112", RUN+="/bin/raw /dev/raw/raw7 %M %m"
ACTION=="add", ENV{MAJOR}=="8", ENV{MINOR}=="128", RUN+="/bin/raw /dev/raw/raw8 %M %m"
ACTION=="add", ENV{MAJOR}=="8", ENV{MINOR}=="144", RUN+="/bin/raw /dev/raw/raw9 %M %m"
ACTION=="add", ENV{MAJOR}=="8", ENV{MINOR}=="160", RUN+="/bin/raw /dev/raw/raw10 %M %m"
ACTION=="add", ENV{MAJOR}=="8", ENV{MINOR}=="192", RUN+="/bin/raw /dev/raw/raw11 %M %m"
ACTION=="add", ENV{MAJOR}=="8", ENV{MINOR}=="208", RUN+="/bin/raw /dev/raw/raw12 %M %m"
ACTION=="add", ENV{MAJOR}=="8", ENV{MINOR}=="224", RUN+="/bin/raw /dev/raw/raw13 %M %m"
ACTION=="add", ENV{MAJOR}=="8", ENV{MINOR}=="240", RUN+="/bin/raw /dev/raw/raw14 %M %m"
# LOCAL IQ_SYSTEM_TEMP NODE 1 - will not be shared but is mounted from the same group -
only rw from node 1
ACTION=="add", ENV{MAJOR}=="65", ENV{MINOR}=="0", RUN+="/bin/raw /dev/raw/raw15 %M %m"
ACTION=="add", ENV{MAJOR}=="65", ENV{MINOR}=="16", RUN+="/bin/raw /dev/raw/raw16 %M %m"
# LOCAL IQ_SYSTEM_TEMP NODE 2 - will not be shared but is mounted from the same group -
only rw from node 2
ACTION=="add", ENV{MAJOR}=="65", ENV{MINOR}=="32", RUN+="/bin/raw /dev/raw/raw17 %M %m"
ACTION=="add", ENV{MAJOR}=="65", ENV{MINOR}=="48", RUN+="/bin/raw /dev/raw/raw18 %M %m"
# IQ_SHARED_TEMP shared rw dbspace file devices
ACTION=="add", ENV{MAJOR}=="65", ENV{MINOR}=="64", RUN+="/bin/raw /dev/raw/raw19 %M %m"
ACTION=="add", ENV{MAJOR}=="65", ENV{MINOR}=="80", RUN+="/bin/raw /dev/raw/raw20 %M %m"
```

```
# SOURCE DATA 400 GB device - not shared and not raw but has a position in the block
devices that will take a slot in the block devices list
# documented here for clarity
# /dev/sdk - source data on filesystem
```

/etc/udev/rules.d/99-raw.rules

```
#/etc/udev/rules.d/99-raw.rules
#change ownership of raw devices to 'sybase'
#change access of /dev/raw/rawctl to allow 'sybase' to read/write
KERNEL=="raw1", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw2", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw3", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw4", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw5", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw6", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw7", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw8", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw9", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw10", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw11", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw12", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw13", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw14", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw15", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw16", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw17", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw18", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw19", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="raw20", OWNER="sybase", GROUP="sybase", MODE="640"
KERNEL=="rawctl", MODE:="666"
```

Resources

Customers can find more information about the VMware and SAP Sybase IQ products described in this paper using the links listed below.

VMware

Deployment References

- VMware vSphere® with Operations Management™:
  http://www.vmware.com/products/datacenter-virtualization/vsphere/overview.html
- VMware vSphere Product Support Center web site:
  http://www.vmware.com/support/product-support/vsphere/index.html
- Performance Best Practices for VMware vSphere 5.1:
  http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.1.pdf
- VMware Cloud Infrastructure Validated by SAP for Virtualization of SAP's Database Portfolio:
  http://www.vmware.com/company/news/releases/vmw-sap-database-05-14-12.html
- VMware vSphere Distributed Switch Best Practices:
  http://www.vmware.com/resources/techresources/10250
- Compatibility Guides (VMware Certified Compatibility Guides):
  http://www.vmware.com/resources/guides.html?src=WWW_BestMatch_US#utm_source=WWW_BestMatch_US&utm_medium=src&utm_campaign=src-tagged-url
- VMware Compatibility Guide (compatible storage and networking devices):
  http://www.vmware.com/resources/compatibility/search.php
- Performance Characterization of VMFS and RDM Using a SAN:
  http://www.vmware.com/files/pdf/performance_char_vmfs_rdm.pdf
- Configuring disks to use VMware Paravirtual SCSI (PVSCSI) adapters (1010398):
  http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1010398
- Disabling simultaneous write protection provided by VMFS using the Multiwriter flag (1034165) (to enable Multiwriter):
  http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1034165
- Change the SCSI Bus Sharing Configuration (to enable SCSI Bus Sharing):
  http://pubs.vmware.com/vsphere-4-esx-vcenter/index.jsp?topic=/com.vmware.vsphere.vmadmin.doc_41/vsp_vm_guide/configuring_virtual_machines/t_change_the_scsi_controller_or_scsi_bus_sharing_configuration.html

General Information

- VMware web site:

  http://www.vmware.com/

- Server & Datacenter Virtualization Products:

  http://www.vmware.com/products/datacenter-virtualization/

- Featured VMware Documentation Sets:
  http://www.vmware.com/support/pubs/

- VMware Licensing Help Center:
  http://www.vmware.com/support/licensing/

VMware Community and TV

- VMware Community, VMware Technology Network (VMTN):
  https://communities.vmware.com/community/vmtn

- VMware Best Practices (see VIOPS links):
  https://communities.vmware.com/community/vmtn/bestpractices

- VMware Community, VMware Knowledge Base:
  http://communities.vmware.com/community/vmtn/resources/knowledgebase

- VMware Support Insider:
  http://blogs.vmware.com/kbtv/

- VMware TV:
  http://www.youtube.com/user/vmwaretv

- VMworld TV:
  http://www.youtube.com/user/VMworldTV

- VMware KBTV (external):
  http://www.youtube.com/user/VMwareKB

SAP Sybase

- SAP Web site:

  http://www.sap.com

- SAP Sybase IQ:
  http://www54.sap.com/pc/tech/database/software/sybase-iq-big-data-management/index.html

- SAP-Sybase Makes Running Business Critical Databases on vSphere Hassle Free:

  http://blogs.vmware.com/alliances/2012/06/vmware_sybase_sap.html

Acknowledgements

These individuals contributed to the creation of this deployment guide:

- **Balayya Kamanboina**, ISV Validation Engineer, VMware
- **Chris Baker**, Solution Advisor Expert, Enterprise Systems Group, SAP Sybase
- **Bob Goldsand**, Staff Partner Architect, VMware
- **Jon Catanzano**, Senior Technical Writer/Editor, Consultant