White Paper

The SGI Adaptive Data Warehouse Solution
An Oracle Optimized Warehouse for the SGI® Altix® XE Server Platform

Table of Contents

1.0 Introduction

The Oracle® Optimized Warehouse Initiative (OWI) enables implementation of optimized data warehousing solutions that combine the high performance and scalability of the most popular database for data warehousing – Oracle® Database 10g – with hardware platforms and storage from industry-leading manufacturers.

OWI enables you to quickly and easily implement Oracle data warehouse solutions. The initiative includes Reference Configurations and Oracle Optimized Warehouses built using industry standard components from multiple hardware platform partners:

– Oracle Optimized Warehouses provide pre-configured and easy-to-deploy data warehouse solutions.
– Oracle Optimized Warehouse Reference Configurations provide the greatest degree of flexibility and choice as you determine your data warehouse platform.

2.0 Oracle Optimized Warehouses

Oracle Optimized Warehouses deliver pre-installed and pre-configured data warehouse solutions. Oracle jointly engineers these configurations with its platform partners providing a higher level of quality assurance and performance. The partnerships marry the enterprise-class capabilities of Oracle Database 10g Enterprise Edition and Oracle® Real Application Clusters with validated hardware configurations optimized for data warehousing. Oracle Optimized Warehouses are jointly released, marketed and sold as a single solution. They are designed for rapid implementation and can quickly provide you with a highly available, scalable, and secure data warehousing solution. You can continue to leverage your existing IT resources and skill sets to support Oracle Optimized Warehouses.

Oracle defines three broad performance categories for Oracle Optimized Warehouse optimized configurations: performance, mainstream, and capacity. Performance systems are optimized for the highest throughput for applications that have the most stringent performance requirements. Mainstream systems address a broad category of users and concurrent workloads. Capacity systems are optimized for environments where storage capacity is of high importance, but where queries do not always access a high percentage of stored data.

Oracle Optimized Warehouse Reference Configurations

Oracle Optimized Warehouse reference configurations provide additional configurations that are based on the validated configurations. The SGI reference configurations listed in Table 2 provide both performance and mainstream implementations that range from 1 TByte to 25 TBytes of user data capacity. With capacity-oriented systems it is possible to extend these configurations to much larger sizes with alternative storage devices, such as 1TB SATA drives which increase capacity by up to 6 times.

3.0 Oracle Optimized Warehouse based on SGI Adaptive Data Warehouse Solution with SGI® Altix® XE250 Server and SGI® InfiniteStorage 220

SGI® Altix XE250 Servers

Data warehouse solutions that utilize SGI® Altix® servers can provide exceptional performance and throughput, and offer substantial headroom for future application growth. The SGI Altix XE250 server is a powerful, compact system that is ideal for scale-out architectures and incremental expansion strategies. Leveraging the strength and versatility of quad-core Intel® Xeon® processors, Altix XE250 servers provide organizations with a fast, reliable platform at an exceptional price point. The innovative board design of SGI Altix XE250 servers maximizes compute density, helping to reduce datacenter floor space and energy expenses.

SGI® InfiniteStorage 220 RAID Products

The SGI InfiniteStorage 220 RAID system offers enterprises cost-effective configurations for low initial capacity requirements and easily expands to support up to 24 TB of total storage. By incorporating an appliance-like design, the InfiniteStorage 220 system simplifies initial deployment and on-going management, and helps organizations to reduce administration burden and improve operational efficiencies. An intuitive GUI helps administrators automate and streamline configuration tasks, such as capacity scaling and RAID group expansion.

4.0 Deployment

SGI configures and integrates the optimized SGI® Adaptive Data Warehouse configurations before shipping. On arrival, SGI field services will perform the physical installation of SGI hardware and connect it into a customer networking environment. As part of the delivery, process SGI field services will perform the system operations diagnosis to make sure all hardware and software components are fully operational before turning the system over to the customer.

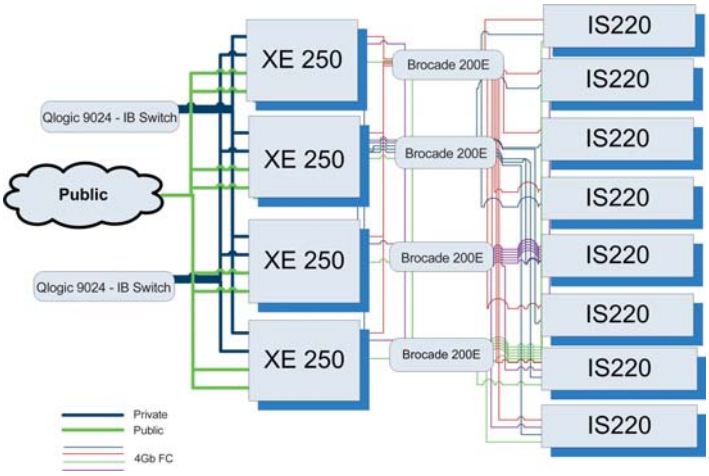

Figure 1: The SGI XE5P Oracle Optimized Warehouse
(Four blocks shown, each block is one XE250 and two IS220 systems)

5.0 Test Methodology

The tested Warehouse is based on a retail schema. The workload is driven by a series of scripts and does not require a third party load generation tool. The schema has 12 tables and is focused around two fact tables, sales, and returns, with the remaining 10 tables acting as dimension tables. The two main fact tables are range partitioned by date and sub-partitioned by hash on their join key. The testing consisted of multiple scenarios. The first scenario consisted of gradually increasing the number of concurrent users executing queries against a single Oracle Optimized Warehouse block. Subsequent tests introduced additional blocks up to a maximum of four. With each addition of a block, the environment was tested by gradually increasing the number of concurrent users against the aggregate Data Warehouse size which grew in units of 1.25TB of user data per block.

As shown in Figure 2 below, user scalability was linear with additional blocks.

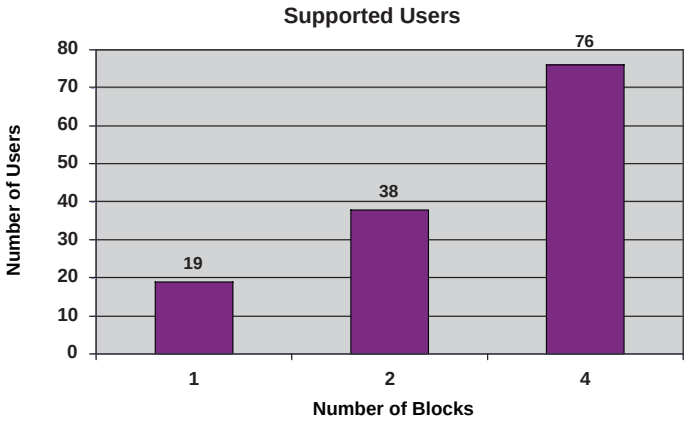

Figure 2: SGI User Scalability

Minimum Hardware/Software Requirements

Minimum Hardware/Software Requirements			
	Validated Components	**Minimum Single Building Block Configuration**	**Minimum Multiple Building Block Configuration**
XE250 Nodes	SGI XE 250	One XE250	Two or more XE250
Memory	All valid memory configurations	32 GB (per node)	32 GB (per node)
SGI Infinite Storage	IS220	2 IS220, with 48 disks each (146 GB 15K SAS)	4 or more IS220, with 48 disks each (146 GB 15K SAS)
Fiber Channel Switches	Brocade 200E		2 switches
HBAs	SGI - FC-10	Dual port 4GB Fibre, 2 (per node)	Dual port 4GB Fibre, 2 (per node)
Ethernet Ports	SGI	4 (per node)	4 (per node)
IB Switches (for Private Interconnect)	Qlogic 9024-CU24-ST2-DDR	2 switches	2 switches
Internal Drives	300 GB SAS	2 (per node) RAID 10	2 (per node) RAID 10
Oracle Software & Licenses	Oracle 10g R2 10.2.0.4 Enterprise Edition, Oracle RAC 10.2.0.4 Oracle ASM 10.2.0.4 Note: ASM Lib is not used		
Operating System	RH5.1/OEL 5.1		

Table I. Hardware requirements for the Oracle Optimized Warehouse on SGI XE250

System Reference Number	1	2	3	4	5	6	7	8	9	10	11	12
	XE1P	XE2.5P	XE5P	XE10P	XE15P	XE20P	XE25P	XE5M	XE10M	XE15M	XE20M	XE25M
Max raw (atomic) user data (TB)	1.56	3.11	6.23	12.46	18.69	24.92	31.15	6.40	12.80	19.20	25.60	32.00
Architecture (Cluster/SMP)	Cluster	Cluster	Cluster	Cluster	Cluster	Cluster	Cluster	Cluster	Cluster	Cluster	Cluster	Cluster
Server Hardware Details												
Server Name	Altix XE	Altix XE	Altix XE	Altix XE	Altix XE	Altix XE	Altix XE	Altix XE	Altix XE	Altix XE	Altix XE	Altix XE
Server Model	XE250	XE250	XE250	XE250	XE250	XE250	XE250	XE250	XE250	XE250	XE250	XE250
# Nodes	1	2	4	8	12	16	20	2	4	6	8	10
CPU Type	Intel Xeon 64bit	Intel Xeon 64bit	Intel Xeon 64bit	Intel Xeon 64bit	Intel Xeon 64bit	Intel Xeon 64bit	Intel Xeon 64bit	Intel Xeon 64bit	Intel Xeon 64bit	Intel Xeon 64bit	Intel Xeon 64bit	Intel Xeon 64bit
# CPUs	2	2	2	2	2	2	2	2	2	2	2	2
# Cores per CPU	4	4	4	4	4	4	4	4	4	4	4	4
Memory per node (server)	32 GB	32 GB	32 GB	32 GB	32 GB	32 GB	32 GB	32 GB	32 GB	32 GB	32 GB	32 GB
Node Interconnect	IB	IB	IB	IB	IB	IB	IB	IB	IB	IB	IB	IB
Storage Hardware												
Storage Disk Array Name / Model	IS220	IS220	IS220	IS220	IS220	IS220	IS220	IS220	IS220	IS220	IS220	IS220
# Storage Arrays	2	4	8	16	24	32	40	4	8	12	16	20
# Disk Controllers per array	2	2	2	2	2	2	2	2	2	2	2	2
# Ports per Disk Controller	2	2	2	2	2	2	2	2	2	2	2	2
Controller I/O Throughput per port (Gb/s)	4	4	4	4	4	4	4	4	4	4	4	4
Disk Transfer Speed	15K RPM	15K RPM	15K RPM	15K RPM	15K RPM	15K RPM	15K RPM	15K RPM	15K RPM	15K RPM	15K RPM	15K RPM
# Total Disks	72	144	288	576	864	1152	1440	144	288	432	576	720
Disk Capacity (GB)	146	146	146	146	146	146	146	300	300	300	300	300
DB Usable Storage (TB)	4.67	9.34	18.69	37.38	56.06	74.75	93.44	19.20	38.40	57.60	76.80	96.00
DW user data (TB)	1.56	3.11	6.23	12.46	18.69	24.92	31.15	6.40	12.80	19.20	25.60	32.00

Table 2. Oracle Optimized Warehouse SGI XE250 Reference Configurations

Hardware Details	1 Block	2 Blocks	4 Blocks
High Performance DW	1 TB	2.5TB	5 TB
Mainstream DW	2.5 TB	5 TB	10 TB
Server Name	Altix	Altix	Altix
Server Model	XE250	XE250	XE250
# Nodes	1	2	4
CPU Type	Intel Xeon 64bit	Intel Xeon 64bit	Intel Xeon 64bit
# CPUs	2	4	8
# Cores per CPU	4 (Quad)	4 (Quad)	4 (Quad)
Memory per node (server)	32 GB	32 GB	32 GB
# Fiber Channel Cards / Node	2 – dual port	2 – dual port	2 – dual port
Node Interconnect	IB	IB	IB
SAN Switch Name		Brocade 200E	Brocade 200E
# Switches		2	4
# Ports per switch	16 ports	16 ports	16 ports
Storage Disk Array Name / Model	IS220	IS220	IS220
# Disk Array Enclosures	2	4	8
# Disk Controllers	2 per array	2 per array	2 per array
# Disks	48 per array	48 per array	48 per array
# disk for Data	32 per array 8 Luns Raid (10) 2+2	32 per array 8 Luns Raid (10) 2+2	32 per array 8 Luns Raid (10) 2+2
# disk for FRA	8 per array 2 Luns Raid (10) 2+2	8 per array 2 Luns Raid (10) 2+2	8 per array 2 Luns Raid (10) 2+2
Disk Transfer Speed	15K	15K	15K
Disk Type	SAS	SAS	SAS
Disk Capacity - RAW Storage TB	146 GB per drive	146 GB per drive	146 GB per drive
Total Usable Storage after RAID 10	7 TB	14 TB	28 TB
Host Bus Adapter	4GB FC dual port	4GB FC dual port	4GB FC dual port
# HBA's	2 per node, each with 2 FC ports	2 per node, each with 2 FC ports	2 per node, each with 2 FC ports

Table 3. SGI XE250 Building block configurations and capacities of the Oracle Optimized Warehouse

Corporate Office
1140 E. Arques Avenue
Sunnyvale, CA 94085
(408) 524-1980
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5185.2860
Europe +44 118.912.7500
Japan +81 3.5488.1811
Asia Pacific +61 2.9448.1463

4123 [11.10.2008] **J15381**