



Configuring and Performance Tuning EMC Symmetrix for Sybase IQ- Multiplex

A technical reference to configure EMC Symmetrix and Sybase IQ-Multiplex in RAID-1 and RAID-S environments for performance optimization.

**EMC Engineering – Solutions and
Software Partner Integration**

**Sybase Inc. – Business Intelligence
Division**

P/N LS031601/REV1

EMC Corporation 171 South Street, Hopkinton, MA 01748-9103

Corporate Headquarters:(508) 435-1000, (800) 424-EMC2

Fax: (508) 435-5374, Service: (800) SVC-4EMC

Copyright © 2001 EMC Corporation. All rights reserved.

Printed 4/17/01

No part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written consent of EMC Corporation.

The information contained in this document is subject to change without notice. EMC Corporation assumes no responsibility for any errors that may appear.

All computer software programs, including but not limited to microcode, described in this document are furnished under a license, and may be used or copied only in accordance with the terms of such license. EMC either owns or has the right to license the computer software programs described in this document. EMC Corporation retains all rights, title and interest in the computer software programs.

EMC Corporation makes no warranties, express or implied, by operation of law or otherwise, relating to this document, the products or the computer software programs described herein. EMC CORPORATION DISCLAIMS ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. In no event will EMC Corporation be liable for (a) incidental, indirect, special or consequential damages or (b) any damages whatsoever resulting from the loss of use, data or profits, arising out of this document, even if advised of the possibility of such damages.

Trademark Information

EMC² (the EMC logo), Symmetrix, TimeFinder, Connectrix, Volume Logix, PowerPath, EMC Enterprise Storage Network and EMC Control Center are registered trademarks of EMC Corporation.

PDB, SPARC, SPARCcenter, SPARCcluster, and SPARCserver, and UltraSPARC are trademarks of SPARC International, Inc. licensed exclusively to Sun Microsystems.

DiskSuite, Network File System, NFS, Solaris, Solstice, Sun, Sun Microsystems and SunOS are trademarks of Sun Microsystems, Inc.

Sybase Adaptive Server IQ with Multiplex and Sybase Central Java Edition are trademarks and Sybase is a registered trademark of Sybase, Inc.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd.

Windows NT is a registered trademark of Microsoft Corporation in the USA and other countries.

All other trademarks used herein are the property of their respective owners.

Table of Contents

| | |
|--|-----------|
| Introduction..... | 1 |
| <i>Overview of Connectrix.....</i> | <i>1</i> |
| <i>Overview of Volume Logix.....</i> | <i>1</i> |
| <i>Overview of PowerPath.....</i> | <i>3</i> |
| <i>Overview of RAID-1 vs RAID-S.....</i> | <i>4</i> |
| RAID-1 | 4 |
| RAID-S | 4 |
| <i>Adaptive Server IQ with Multiplex Overview.....</i> | <i>5</i> |
| <i>Project Configurations</i> | <i>5</i> |
| System Configurations | 5 |
| Configuring the Test Environment..... | 6 |
| Application Installation and Configuration..... | 12 |
| <i>IQ-Multiplex Cache.....</i> | <i>12</i> |
| <i>IQ-Multiplex Indexing.....</i> | <i>12</i> |
| Hardware Testing | 14 |
| <i>RAID-1 Hardware Tests</i> | <i>15</i> |
| <i>RAID-S Hardware Tests</i> | <i>16</i> |
| <i>Observations</i> | <i>16</i> |
| <i>Hardware Test Summary.....</i> | <i>17</i> |
| Application Testing | 18 |
| <i>Descriptions of Tables.....</i> | <i>18</i> |
| <i>Data I/O Rates</i> | <i>20</i> |
| <i>Sorts</i> | <i>20</i> |
| <i>Index Performance.....</i> | <i>21</i> |
| <i>Application Test Summary.....</i> | <i>22</i> |
| Additional Documentation..... | 23 |

Introduction

This paper documents the configuration and best practices for performance tuning EMC Symmetrix® and Sybase® Adaptive Server® IQ with Multiplex (herein referred to as IQ-Multiplex) in both RAID-1 and RAID-S environments. The purpose of this white paper is to:

- Provide a general understanding of how to configure EMC Symmetrix and Sybase IQ-Multiplex in conjunction with Connectrix™, PowerPath™, Fibre and Volume Logix™.
- Provide guidelines and best practice procedures for configuring metavolumes in different RAID environments to optimize performance.

The procedures and guidelines in this paper were established as a result of lab testing within EMC Engineering. The hardware, software, and installation configurations that were used during the testing are also described in this paper.

The first set of tests were performed to test the hardware configuration from the front end (hosts) all the way to the back end (disks). Our intent was to ensure that the configuration of metavolumes, Fibre Channel Director/Interface Adapters (FAs), Symmetrix cache, and Disk Adapters (DAs) would support the data warehouse application after n terabytes were loaded. The last (second) set of tests were performed to test the application in its respective EMC Symmetrix environment. The purpose of this document is to serve as a baseline configuration test plan for customers planning to install Sybase IQ-Multiplex on EMC Symmetrix.

Overview of Connectrix

The EMC Connectrix Enterprise Storage Network System™ delivers significant enhancements to Symmetrix open systems connectivity. The Symmetrix-to-server connectivity architecture provided by Connectrix is based on the Fibre Channel fabric topology (FC-SW). Fabric support in the Connectrix is built on high-speed optical and digital switching technologies. Enterprise-class redundancy features are built on proven technologies used in ESCON connectivity. The protocols used in the Connectrix are based on industry standards. Fibre Channel fabric technology is evolving rapidly. The release of Symmetrix-based solutions mirror the practices used today for qualifying new servers, clustering technologies, operating systems, drivers, and host bus adapters. For more detailed information on EMC Connectrix refer to the following web page:

http://www.emc.com/products/networking/connectrix.jsp?openfolder=storage_networkingVolume_Logix_Overview

Overview of Volume Logix

In a typical Fibre Channel topology, each host is connected to the Fibre Channel hub or switch through one or more host bus adapter (HBA) ports. Each Symmetrix unit is connected to the Fibre Channel hub or switch through one or more Fibre Channel Director/Interface Adapter (FA) ports, each FA port providing access to a given set of volumes. Each port in the Fibre Channel topology (HBA and Symmetrix FA) is uniquely identified by its World Wide Name (WWN). Volume Logix is volume configuration management (VCM) software that controls access to Symmetrix volumes. Unlike host-based access control mechanisms that are error prone (since each host must be trusted not to conflict with others), the Symmetrix-based access control mechanism provided by Volume Logix avoids such conflicts by single centralized monitoring.

Using Volume Logix, you can define a *virtual channel* connecting each host with its storage volumes in the Symmetrix unit, even though there may be many hosts sharing the same Symmetrix port. When a host attempts to access a storage volume on the Symmetrix unit, the host HBA's WWN, as supplied during login to the Fabricvolumeor Arbitrated Loop, is passed to the Symmetrix FA port. The Symmetrix volume records the connection, stores the WWN in a *Login History table*, and builds a filter listing the volumes available to that WWN through that FA port. The HBA port then sends I/O requests directed at a particular Symmetrix volume to the FA port. Each request includes the identity of the requesting HBA (from which its WWN can be determined) and the identity of the requested volume, with its FA and logical unit number (LUN). The Volume Logix software that runs on a Symmetrix unit processes each I/O request to verify that the HBA is allowed to access that volume. Any request for a volume that an HBA does not have access to returns an error to the host.

When the Symmetrix unit is configured, a special device called the VCM device must be created. The size of the Symmetrix device designated as the Volume Logix database device is 16 Symmetrix cylinders (approximately 7.5 MB). All space on this device is reserved for current and future Volume Logix functions. There is one Volume Logix database for each Symmetrix unit and it should be locally mirrored. The Volume Logix database contains records that specify the devices that may be accessed by a particular WWN through a specific FA. Each FA may control the access for as many as 32 unique WWNs. As many as 16 FA ports (depending on the Symmetrix type) may be configured within the Volume Logix database.

After the Volume Logix software is installed, use the `vcmfind` utility on each host attached to the Symmetrix unit in your environment. The `vcmfind` command does the following:

- determines the WWN of each host HBA
- assigns a recognizable alias known as an ASCII WWN (AWWN), and
- enables the Graphical User Interface (GUI) to draw and display the topology. If an AWWN is not present, the GUI will not be able to create an accurate picture of the EMC Enterprise Storage Network™ (ESN).

Figure 1 shows a sample representation of the Volume Logix GUI. The left half of the screen shows several hosts with their respective adapters (HBAs), the fibre port numbers, and the Symmetrix units. The right half of the screen shows the individual devices; that is, the logical volumes, metavolumes, business continuance volumes (BCVs) and the VCM device. The devices underlined in black indicate that the device has been assigned to a host. The device surrounded by a blue box indicates that it has been assigned to the host that is highlighted.

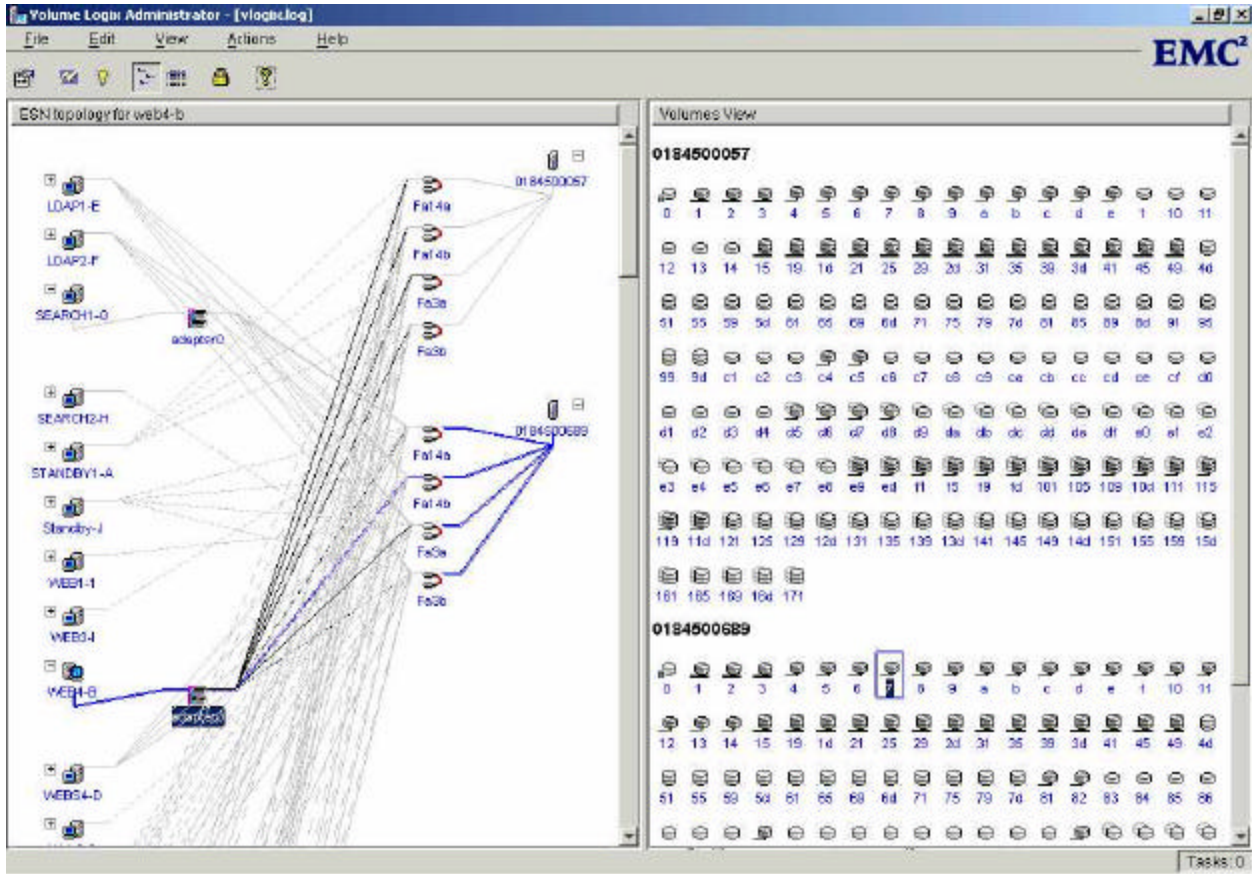


Figure 1. Sample Representation of the Volume Logix GUI

Overview of PowerPath

EMC PowerPath is host-based software that works with Symmetrix. PowerPath enhances Symmetrix capabilities by providing automatic detection of and recovery from host-to-Symmetrix connection failures. PowerPath enables you to establish from two to 32 SCSI paths or 16 Fibre Channel paths between a host and a Symmetrix device configured for the host. You can use these paths simultaneously for data I/O to the device. If a failure occurs on any path, PowerPath automatically redirects data over the other paths without disrupting I/O. In addition, PowerPath dynamically balances the distribution of the I/O load to improve performance. You can configure data paths, maintain the paths, and tune for performance while on line.

PowerPath offers the following features and benefits:

- **Multipath support** — provides multiple channel access (up to 32) to a Symmetrix device.
- **Load balancing** — automatically adjusts data routing for optimum performance, and eliminates the need to statically configure devices among multiple channels.

- **Path failover** — automatically and nondisruptively redirects data to an alternate data path; eliminates application downtime in the event of path failure.
- **Online recovery** — allows you to resume use of a path after the path is repaired without service interruption.

For more information on EMC PowerPath, refer to the following web page:

http://www.emc.com/products/software/PowerPath.jsp?openfolder=storage_software

Overview of RAID-1 vs RAID-S

This section provides a brief overview of the functionality of RAID-1 and RAID-S disk configurations.

RAID-1

RAID-1 mirroring provides the highest level of performance and availability for all mission-critical and business-critical applications. RAID-1 mirroring maintains a duplicate copy of each logical volume on two physical disk devices. Symmetrix maintains these copies internally, transparent to the host(s), by writing all modified data to both devices. Symmetrix designates two logical volumes residing on different physical disks as a mirrored pair — one volume is mirror-1 (M1) and the other volume is mirror-2 (M2). The host(s) view the M1 and M2 volumes as the same logical volume because each has the same address.

To ensure the highest availability, each volume is attached to separate Disk Directors that are attached to different memory buses. The EMC implementation of RAID-1 mirroring on Symmetrix systems includes performance enhancements beyond the high-availability capabilities normally associated with RAID-1. A write operation to any mirrored volume is executed identically to a non-mirrored write. The Channel Director presents Channel End/Device End to the host after data is written and verified in cache. The Disk Directors then asynchronously destage the data to each drive of the mirrored pair of drives.

As such, mirroring on Symmetrix systems exploits the 100% fast write capability, and the application does not incur additional time associated with having to physically perform two disk write I/Os (one to each drive of the mirrored pair) that is normally associated with RAID-1. The Symmetrix performance algorithms for read operations in mirrored pairs offer three mirroring policies to best balance the use of the Symmetrix architecture: interleave, split, and dynamic mirrored service policy (DMSP). For more details and information related to RAID-1 and/or RAID-S refer to the following web page:

http://www.emc.com/products/product_pdfs/pdg/symm8000_c756.pdf

RAID-S

RAID-S for the Symmetrix system provides parity-based data protection similar to RAID-4 and RAID-5, but with some significant advantages for performance, flexibility, and data availability. Compared to a RAID-1 mirrored Symmetrix, RAID-S offers more usable capacity than a mirrored system containing the same number of disk drives. Like mirroring, RAID-S parity protection can be dynamically added or removed. For example, for higher performance requirements and high availability, parity protection on a RAID-S group can be turned off and the volumes in the RAID-S group mirrored. Within the same Symmetrix system, data can be protected via RAID-S, mirroring, and/or the Symmetrix remote data facility™ (SRDF). Dynamic sparing can be added to any of these data protection options.

One of the factors contributing to the higher performance of the Symmetrix RAID-S option is that Symmetrix takes advantage of the ability of the latest disk drives to calculate parity at the disk itself. The Boolean operation “Exclusive Or” (XOR) logic used to calculate the parity is carried out by a microprocessor with XOR logic and disk cache on each disk drive. This greatly improves write performance by offloading these calculations from the host or Symmetrix system, allowing them to continue to service I/O requests. Since the Channel Directors do not need to calculate parity and I/Os are

serviced from cache, RAID-S will not impose performance penalties on the host processor. EMC currently recommends that data protected by RAID-S be *grouped* with a ratio of three data disks to one parity disk.

Though Symmetrix 8000 allows the use of different capacity disks within a single Symmetrix system, all physical disks participating in a RAID-S group must have identical storage capacity. With this approach, 75% of the total storage capacity of each RAID-S group of volumes is available for storing data. Multiple RAID-S groups may exist within a single Symmetrix system. Members of a RAID-S group can be located anywhere in the Symmetrix system, spanning multiple Disk Directors. Currently EMC does not support a RAID-S configuration on disk drives greater than 50 GB. For more details and information related to RAID-1and/or RAID-S, refer to the following web page:

http://www.emc.com/products/product_pdfs/pdg/symm8000_c756.pdf

Adaptive Server IQ with Multiplex Overview

IQ-Multiplex is a high-performance decision support server designed specifically for data warehousing. This cross-platform product runs on Windows NT® as well as on Sun Solaris™ (SPARC™), HP 9000/800 HP-UX, IBM RISC System/6000 AIX, Silicon Graphics IRIX, and Compaq Tru64 UNIX® systems.

IQ-Multiplex is part of the Adaptive Server product family that includes Adaptive Server Enterprise for enterprise transaction and mixed workload environments, and Adaptive Server Anywhere, a small footprint version of Adaptive Server often used for mobile and occasionally connected computing.

The IQ-Multiplex product was qualified with EMC TimeFinder in Q3 2000. For more information on Sybase IQ-Multiplex and IQ-Multiplex Architecture, refer to the white paper *Scalability and Reliability with EMC Symmetrix and Sybase IQ-Multiplex*.

Project Configurations

The operating system, hardware, and test-environment configurations used in testing for this project are documented on the following pages. Note that performance will be directly impacted (positively or negatively) by the configuration of the environment as a whole. Our intent was to create an environment purely designed to push as much I/O bandwidth as possible from the hosts to the Symmetrix system. The project configurations, hardware, and application (IQ-Multiplex) test results and conclusions are documented in the remainder of this paper.

System Configurations

Tables 1 and 2 describe the hardware and software configurations that were used for testing.


Table 1. Hardware Configurations

| Hardware | Details |
|------------------------------|--|
| Sun Ultra Enterprise 4500 | Database (IQ-Multiplex write server); four 336 Mhz CPUs; 1.7 GB Memory |
| Sun Ultra 2 | Database (IQ-Multiplex query server); two 400 Mhz CPUs; 2 GB Memory |
| Jaycor Fibre Channel Adapter | FCA driver v2.2.4.EMC |
| EMC Symmetrix | 8430, 96 - 36 GB drives, microcode 5566, 8 GB cache |
| EMC Symmetrix | 8430, 96 - 73 GB drives, microcode 5566, 8 GB cache |
| EMC Connectrix | ED -1032, Connectrix Manager 2.1 |

Table 2. Software Configurations

| Software | Version |
|--|---------------------------|
| Operating system | Sun Solaris 5.7 64-bit |
| EMC Solutions Enabler (SYMCLI) | 4.1-131 rev. 7 |
| EMC Volume Logix | 2.2.1 |
| EMC Control Center (Workload analyzer) | 4.1 |
| Database Management System | Sybase IQ 12.4.2 (64-bit) |
| PowerPath | 2.0 |

Figure 2 illustrates the server and Symmetrix configuration used for testing.

**Figure 2. Configuration of the Server and the Symmetrix System**

Configuring the Test Environment

After analyzing various configuration methods with the EMC performance team, we created bin files for both the RAID-1 and RAID-S Symmetrix units. We were limited to creating metavolumes no greater than 128 GB, as IQ-Multiplex currently has a single device size limitation.

Configuration of Symmetrix unit and the entire SAN was driven primarily by delivering as much I/O bandwidth as possible to the hosts. The Symmetrix unit and its bandwidth is the only shared component in IQ-Multiplex. The total capacity for bandwidth in the IQ-Multiplex server will determine the number of nodes (hosts) that will achieve 100% scalability. IQ-Multiplex storage management is effective yet simple: any number of hardware devices (preferably raw devices) are mapped to IQ-Multiplex *dbspaces* that

create a storage pool that is managed by the IQ-Multiplex storage manager. The *dbspaces* do not have to be the same size. However, for optimal performance all devices should be configured the same to achieve maximum I/O performance.

Tables are automatically mapped across all *dbspaces* without the intervention of a database administrator (DBA). Experience with many installations shows that the effective "rule of thumb" in sizing I/O bandwidth of the shared disk is 10 MB per second per CPU for IQ-Multiplex.

For example, an IQ-Multiplex with 16 CPUs would need 160 MB/sec (16 CPUs x 10 MB/sec) to push the maximum I/O through each CPU and eliminate I/O wait time during query executions.

The ideal configuration should be able to meet the following goals:

- Perform automatic I/O load balancing across all *dbspaces*.
- Eliminate *hot spots* resulting from uneven data distribution (*data skew*).
- Eliminate *hot spots* resulting from uneven data access (*query skew*).
- Simplify DBA task of configuring and optimizing the IQ-Multiplex.

Configuration and Performance Tips

Configuration: the use of large (100+ GB) raw devices reduces the total number of devices needed (For example: 8 x 125 GB devices = 1 TB database).

Performance optimization: after the Symmetrix unit, hosts, and IQ-Multiplex have been configured, the system will automatically load balance across all devices; consume device space evenly, and deliver optimal bandwidth to all hosts in the multiplex.

Striping was implemented on both the Symmetrix unit and within the IQ-Multiplex environment. Our intent was to create an environment designed to push as much I/O bandwidth as possible from the hosts to the Symmetrix units. IQ-Multiplex is designed for ad hoc queries of any size and complexity that can access a small portion of the data or the entire data set. Due to the IQ-Multiplex shared disk non-partitioned architecture, the main data store must be designed to deliver maximum data flow (bandwidth) to nodes in the IQ-Multiplex environment and to eliminate *hot spots* (or areas on disk that are frequently accessed and may affect overall performance).

To eliminate *hot spots* on a Symmetrix system, do the following:

- Stripe each metavolume using Symmetrix hardware striping for maximum bandwidth. RAID-1 and RAID-S striping details are explained later in this paper. IQ-Multiplex uses large I/O sizes (64 KB to 512 KB) and metavolumes must be striped accordingly.

To eliminate *hot spots* with IQ-Multiplex, do the following:

- Use IQ-Multiplex striping that automatically stripes IQ pages across all *dbspaces* in the database. IQM striping provides extremely efficient I/O load balancing across all *dbspaces* that are mapped to metavolumes. This eliminates any need for partitioning and placement control and consequently simplifies very large data base (VLDB) design and administration, and eliminates the need for I/O tuning.

This combination of hardware striping and IQ-Multiplex striping delivers excellent I/O performance and maximizes all available bandwidth.

The ability of IQ-Multiplex to perform prefetch algorithms in combination with the EMC Symmetrix caching algorithms worked very well together.

On the Symmetrix unit configured for RAID-1, three different environments were set up to accommodate different striping configurations. The final input for bin file creation consisted of the following scenarios:

- Config 1 consists of fourteen 8.42-GB volumes, for a total metavolume size of 117.8 GB. All metavolumes in Config 1 have a 2-cylinder stripe (or 916 KB stripe). A 2-cylinder stripe is the

minimum allowable stripe size for a Symmetrix metavolume. Eight metavolumes were created for each group of configured disks (called a config group), totaling 943 GB for this config group.

- Config 2 consists of fourteen 8.42-GB volumes, for a total metavolume size of 117.8 GB. All metavolumes in Config 2 have an 8-cylinder stripe (or 3.7 MB stripe). Eight metavolumes were created for this config group, totaling 943 GB.
- Config 3 consists of fourteen 8.42-GB volumes, for a total metavolume size of 117.8 GB. All metavolumes in Config 3 have a 16-cylinder stripe (or 7.3 MB stripe). Eight metavolumes were created for this config group, totaling 943 GB.

The total disk configured for the three RAID-1 environments is 2,829 GB (2.8 TB).

On the Symmetrix unit configured for RAID-S, three different environments were set up to accommodate different striping configurations. The final input for bin file creation consisted of the following scenarios:

- Config 1 consists of fourteen 8.42-GB volumes, for a total metavolume size of 117.8 GB. All metavolumes in Config 1 have a 2-cylinder stripe (or 916 KB stripe). A two cylinder stripe is the *minimum* allowable stripe size for a Symmetrix metavolume. Unlike the Symmetrix unit configured for RAID-1 which contained 73 GB disk drives, this Symmetrix unit configured for RAID-S had 36 GB disk drives. There was not enough disk space to create eight groups of metavolumes. Instead, six metavolumes were created for each config group, totaling 707 GB per config group.
- Config 2 consists of fourteen 8.42-GB volumes, for a total metavolume size of 117.8 GB. All metavolumes in Config 2 have an eight cylinder stripe (or 3.7 MB stripe). Six metavolumes were created for this config group, totaling 707 GB.
- Config 3 consists of fourteen 8.42-GB volumes, for a total metavolume size of 117.8 GB. All metavolumes in Config 3 have a 16-cylinder stripe (or 7.3 MB stripe). Six metavolumes were created for this config group, totaling 707 GB.

The total disk configured for the three RAID-S environments is 2,121 GB (2.1 TB).

Table 3 summarizes the six different configurations that were created for the testing. Three of the configurations are RAID-1 and three are RAID-S.

Table 3. Summary of the Six Different Testing Configurations

| Configuration | RAID | Logical Device Size | # Devices | # Metas | Stripe Size | PowerPath Enabled |
|---------------|------|---------------------|-----------|---------|---------------------|-------------------|
| 1 | 1 | 8.42 | 14 | 8 | 2 cyl. (916 K) | Yes |
| 2 | 1 | 8.42 | 14 | 8 | 8 cyl. (3.7 MB) | Yes |
| 3 | 1 | 8.42 | 14 | 8 | 16 cyl. (7.3 MB) | Yes |
| 1 | S | 8.42 | 14 | 6 | 2 cyl. (916 K) | Yes |
| 2 | S | 8.42 | 14 | 6 | 8 cyl. (3.7 MB) | Yes |
| 3 | S | 8.42 | 14 | 6 | 16 cyl. (7.3 MB) | Yes |

Table 4 lists the RAID-1 raw devices were used for the IQ data store on both hosts. File system devices were used for the IQ catalog store only. Host-based symbolic links were used to access the devices. "M2" signifies the 2-cylinder stripe devices, "M8" signifies the 8-cylinder stripe devices, and "M16" signifies the 16-cylinder stripe devices.

Table 4. Description of RAID-1 Devices Used for Hardware and Application Testing

| EMC Power Device | Symbolic Link | Symmetrix ID # | Symm Dev # | # FA Ports for Device |
|-----------------------|---------------|----------------|------------|-----------------------|
| dev/rdisk/emcpower32c | M2_1 | 778 | 000 | 2 |
| dev/rdisk/emcpower0c | M2_2 | 778 | 00f | 2 |
| dev/rdisk/emcpower18c | M2_3 | 778 | 01e | 2 |
| dev/rdisk/emcpower19c | M2_4 | 778 | 02d | 2 |
| dev/rdisk/emcpower26c | M2_5 | 778 | 03c | 2 |
| dev/rdisk/emcpower27c | M2_6 | 778 | 04b | 2 |
| dev/rdisk/emcpower39c | M2_7 | 778 | 05a | 2 |
| dev/rdisk/emcpower40c | M2_8 | 778 | 069 | 2 |
| dev/rdisk/emcpower33c | M8_1 | 778 | 078 | 2 |
| dev/rdisk/emcpower34c | M8_2 | 778 | 087 | 2 |
| dev/rdisk/emcpower20c | M8_3 | 778 | 096 | 2 |
| dev/rdisk/emcpower21c | M8_4 | 778 | 0a5 | 2 |
| dev/rdisk/emcpower28c | M8_5 | 778 | 0b4 | 2 |
| dev/rdisk/emcpower29c | M8_6 | 778 | 0c3 | 2 |
| dev/rdisk/emcpower41c | M8_7 | 778 | 0d2 | 2 |
| dev/rdisk/emcpower42c | M8_8 | 778 | 0e1 | 2 |
| dev/rdisk/emcpower35c | M16_1 | 778 | 0f0 | 2 |
| dev/rdisk/emcpower36c | M16_2 | 778 | 0ff | 2 |
| dev/rdisk/emcpower22c | M16_3 | 778 | 10e | 2 |
| dev/rdisk/emcpower23c | M16_4 | 778 | 11d | 2 |
| dev/rdisk/emcpower30c | M16_5 | 778 | 12c | 2 |
| dev/rdisk/emcpower31c | M16_6 | 778 | 13b | 2 |
| dev/rdisk/emcpower43c | M16_7 | 778 | 14a | 2 |
| dev/rdisk/emcpower37c | M16_8 | 778 | 159 | 2 |

In most cases, the raw devices were used on both hosts. Table 5 lists the RAID-S devices used for hardware and application testing. Host-based symbolic links were used to access the devices. "S2" signifies the 2-cylinder stripe devices, "S8" signifies the 8-cylinder stripe devices, and "S16" signifies the 16-cylinder stripe devices.

Table 5. Description of RAID-S Devices Used for Hardware and Application Testing

| EMC Power Device | Symbolic Link | Symmetrix ID # | Symm Dev # | # FA Ports for Device |
|-----------------------|---------------|----------------|------------|-----------------------|
| dev/rdisk/emcpower11c | S2_1 | 791 | 000 | 2 |
| dev/rdisk/emcpower12c | S2_2 | 791 | 00f | 2 |
| dev/rdisk/emcpower1c | S2_3 | 791 | 01e | 2 |
| dev/rdisk/emcpower2c | S2_4 | 791 | 02d | 2 |
| dev/rdisk/emcpower15c | S2_5 | 791 | 03c | 2 |
| dev/rdisk/emcpower16c | S2_6 | 791 | 04b | 2 |
| dev/rdisk/emcpower17c | S8_1 | 791 | 069 | 2 |
| dev/rdisk/emcpower3c | S8_2 | 791 | 078 | 2 |
| dev/rdisk/emcpower4c | S8_3 | 791 | 087 | 2 |
| dev/rdisk/emcpower7c | S8_4 | 791 | 096 | 2 |
| dev/rdisk/emcpower8c | S8_5 | 791 | 0a5 | 2 |
| dev/rdisk/emcpower45c | S8_6 | 791 | 05a | 2 |
| dev/rdisk/emcpower13c | S16_1 | 791 | 0b4 | 2 |
| dev/rdisk/emcpower14c | S16_2 | 791 | 0c3 | 2 |
| dev/rdisk/emcpower5c | S16_3 | 791 | 0d2 | 2 |
| dev/rdisk/emcpower6c | S16_4 | 791 | 0e1 | 2 |
| dev/rdisk/emcpower9c | S16_5 | 791 | 0f0 | 2 |
| dev/rdisk/emcpower10c | S16_6 | 791 | 0ff | 2 |

Volume Logix was used to assign specific devices to an HBA and FA port. All Symmetrix devices were made visible to all Symmetrix FA ports.

Since we had two HBAs in each server, each device was made visible via two Symmetrix Fibre (or FA) ports to each server. This allowed PowerPath to load balance between HBAs, and between the Symmetrix Fibre ports. Both hosts see each device on separate FA ports, so when each host accesses a device, four separate Fibre ports are being used.

We tested three different metavolume stripe sizes: 2 cylinders, 8 cylinders, and 16 cylinders. In our tests, we used six metavolumes with different metavolume stripe sizes. We tested each stripe size separately. To achieve the greatest throughput, two metavolumes from each stripe set were assigned to a pair of Fibre ports. Thus, we tested all volumes of a particular stripe size, accessing all six Symmetrix Fibre ports. This configuration ensured that there would be no I/O bottlenecks on the six Symmetrix Fibre ports. Table 6 represents the assignments of the Symmetrix volumes to FA ports.

Table 6. Symmetrix Volume Assignments to FA Ports

| Symmetrix | Metavolume Stripe Size | Host #1 - Losas160 FA Port Assignments | Host #2 - Losas161 FA Port Assignments |
|------------------|-------------------------------|---|---|
| 778 | 2-Cylinder (using 6 metas) | 3a/14b, 3b/4a, 4b/14a | 4b/14a, 3a/14b, 3b/4a |
| 778 | 8-Cylinder (using 6 metas) | 3a/14b, 3b/4a, 4b/14a | 4b/14a, 3a/14b, 3b/4a |
| 778 | 16-Cylinder (using 6 metas) | 3a/14b, 3b/4a, 4b/14a | 4b/14a, 3a/14b, 3b/4a |
| | | | |
| 791 | 2-Cylinder (using 6 metas) | 4b/14a, 3a/14b, 3b/4a | 3a/14b, 3b/4a, 4b/14a |
| 791 | 8-Cylinder (using 6 metas) | 4b/14a, 3a/14b, 3b/4a | 3a/14b, 3b/4a, 4b/14a |
| 791 | 16-Cylinder (using 6 metas) | 4b/14a, 3a/14b, 3b/4a | 3a/14b, 3b/4a, 4b/14a |

Application Installation and Configuration

IQ-Multiplex 12.4.2 was loaded on both Sun Servers. The Sun Ultra Enterprise 4500 was set up as the writer node for the IQ-Multiplex and the Sun Ultra 2 was set up as a reader node. Standard installation procedures were used. Refer to the IQ-Multiplex documentation for more information related to IQ-Multiplex installation and setup at: <http://sybooks.sybase.com/mxp1240e.html>

The objective of the test was to build a multiple terabyte IQ-Multiplex 12.4.2 database on an EMC Symmetrix system. Data loading speeds were not a major consideration for the testing due to the processing power of the Symmetrix, the number of CPUs, and the amount of memory on the SUN servers.

For the testing, a database with 24 devices (metavolumes) was allocated to the IQ-Multiplex server. IQ-Multiplex technology will automatically stripe across these raw devices. This is discussed in more detail later.

Due to the the intended size of the database, it was determined that a 256 KB page size would be optimal. The default page size in IQ-Multiplex (12.4.2) is 64 KB.

The devices in the Symmetrix unit were divided between the main database space and the temporary space (referred to as *dbspaces*). IQ-Multiplex uses the temp *dbspaces* for data loading (from the write server) and query processing (from the read servers). In this test, we allocated more temporary space than would be needed under normal circumstances. This was to further enhance performance for IQ-Multiplex. Our goal was to allocate and utilize every device that was created for this project.

Two Symmetrix units were configured: one containing 24 RAID-1 (mirrored) metavolumes, each 117.8 GB in size, and one containing 18 RAID-S (3 + 1 parity) metavolumes, each 117.8 GB in size. One IQ-Multiplex database was created primarily on mirrored disk, and another database was created using the RAID-S devices.

IQ-Multiplex Cache

For most tests, the Adaptive Server IQ-Multiplex cache was configured at 800 MB for main cache and 300 MB for temp cache. The original table (named CDR) was created for the tests with 121 columns and 113 LF/HG (Low_Fast/High_Group) indexes. We used this table to determine optimal (IQ-Multiplex) cache configurations for load performance on the very large tables.

IQ-Multiplex Indexing

The test results in the following pages contain references to different types of indexes. IQ indexes are used to improve data retrieval performance. It's important to grasp a basic understanding of IQ indexing structure and index types.

IQ-Multiplex indexes represent and store the data so that it can be used for processing queries. This strategy is designed for the data warehousing environment, in which queries typically examine enormous numbers of records, often with relatively few unique values, and where aggregate results are commonly required.

When data is loaded into a table, IQ stores data by column rather than by row, for each column in each table. The column orientation gives IQ indexes important performance advantages over traditional row-based indexing.

IQ-Multiplex automatically creates a default index, also known as the Fast Projection index, for each column. This index is used in a number of SQL operations including column projection, table joins, and string searches. When a column is designated as either a Primary Key or Unique, a High_Group (HG) index is automatically created. To achieve maximum query performance, it is recommended that one or more additional index types are defined to best represent the cardinality and use of column data.

IQ-Multiplex index types are as follows:

- **Low_Fast (LF)** - A value-based bitmap for processing queries on low-cardinality data (recommended for up to 1,000 distinct values, but can support up to 10,000).
- **High_Group (HG)** - An enhanced b-tree index to process equality and group by operations on high-cardinality data (recommended for more than 1,000 distinct values).
- **High_Non_Group (HNG)** - A non-value-based bitmap index ideal for most high-cardinality DSS operations involving ranges or aggregates.

Hardware Testing

The mirrored and the RAID-S configurations were tested to ensure that both could produce acceptable levels of performance. Two different test programs were used: `iorate` and `dd`. `Iorate` is an EMC-developed benchmark procedure. Source code and compiled binaries are available on the EMC public FTP site at the following web page: <ftp://ftp.emc.com/pub/symm3000/iorate>. `dd` is a Unix-based utility that copies an input file to an output source, and allows the user to apply any desired conversions in the process. Since IQ-Multiplex is primarily a sequential workload, `dd` was also deemed to be an acceptable test of sequential read performance.

Our tests did not approach Symmetrix performance limits. For this test, the host configuration and the number of Fibre Channel adapters was inadequate to fully drive the Symmetrix unit. This is typical of many installations: the Symmetrix unit will provide more aggregate bandwidth than the host can use.

Either `iorate` or `dd` may be used to test new configurations before installing IQ-Multiplex. A quick test of sequential throughput can uncover issues such as missing host configuration parameters or Symmetrix configuration problems.

Figure 3 presents a sample `read` script that was used to test throughput for a single metavolume.

```
#!/bin/ksh
#
# use this script for SINGLE METAVOLUME test ONLY
# DO NOT use it inside other scripts
# use med to build scripts
#
# Arguments:
# $1 is the symbolic link (Symmetrix metavolume) to test
# $2 is the blocksize for DD
# $3 is the offset (for ISEEK)
# $4 is the count (number of blocksize)

echo "START_HOT_READ=$1; blocksize=$2; offset=$3; count=$4"
date
dd if=/bench/links/$1 of=/dev/null bs=$2 isseek=0 count=$4 &
dd if=/bench/links/$1 of=/dev/null bs=$2 isseek=$3 count=$4 &
dd if=/bench/links/$1 of=/dev/null bs=$2 isseek=2x$3 count=$4 &
dd if=/bench/links/$1 of=/dev/null bs=$2 isseek=3x$3 count=$4 &
dd if=/bench/links/$1 of=/dev/null bs=$2 isseek=4x$3 count=$4 &
dd if=/bench/links/$1 of=/dev/null bs=$2 isseek=5x$3 count=$4 &

i=1
numDD=`ps -ef | grep dd | grep -v grep | wc -l`
while [ $numDD != "0" ] ; do
    sleep 1
    i=`expr $i + 1`
    numDD=`ps -ef | grep dd | grep -v grep | wc -l`
done
date

echo "END_HOT_READ=$1; blocksize=$2; offset=$3; count=$4 DURATION=$i sec"
echo "==="
```

Figure 3. Sample Read Script Used in Testing Throughput for a Single Metavolume

Figure 4 is a sample *write* script that was used to test throughput for a single metavolume.

```
#!/bin/ksh
# med does NOT time anything, use it as BUILDING BLOCK
#
# Arguments:
# $1 is the symbolic link (Symmetrix metavolume) to test
# $2 is the blocksize for DD
# $3 is the offset (for ISEEK)
# $4 is the count (number of blocksize)

echo "START_HOT_WRITE test=$1; blocksize=$2; offset=$3; count=$4"
date
dd if=/dev/zero of=/bench/links/$1 bs=$2 iseek=500 count=$4 &
dd if=/dev/zero of=/bench/links/$1 bs=$2 iseek=$3 count=$4 &
dd if=/dev/zero of=/bench/links/$1 bs=$2 iseek=2x$3 count=$4 &
dd if=/dev/zero of=/bench/links/$1 bs=$2 iseek=3x$3 count=$4 &
dd if=/dev/zero of=/bench/links/$1 bs=$2 iseek=4x$3 count=$4 &
dd if=/dev/zero of=/bench/links/$1 bs=$2 iseek=5x$3 count=$4 &

i=1
numDD=`ps -ef | grep dd | grep -v grep | wc -l`
while [ $numDD != "0" ] ; do
    sleep 1
    i=`expr $i + 1`
    numDD=`ps -ef | grep dd | grep -v grep | wc -l`
done
date

echo "END_HOT_WRITE=$1; blocksize=$2; offset=$3; count=$4 DURATION=$i sec"
echo "==="
```

Figure 4. Sample Write Script Used in Testing Throughput for a Single Metavolume

RAID-1 Hardware Tests

The RAID-1 Symmetrix unit was tested both with *iorate* and with *dd* to ensure that sequential bandwidth was adequate. Neither *iorate* or *dd* was able to fully drive the Symmetrix unit, but the tests did show the host Fibre Channel adapters to be the limiting component.

The following read and write tests were performed:

- Metavolume stripe sizes of 2 cylinders, 8 cylinders, and 16 cylinders

EMC only supports 2-cylinder stripe sizes for customers. Larger stripe sizes were tested simply to note any differences in sequential performance.

- I/O sizes 64 KB, 128 KB, 256 KB, and 512 KB for *iorate*
- I/O sizes 64 KB, 128 KB, 256 KB, and 1,024 KB for *dd*
- Single metavolume and six metavolumes
- Single stream per device and multiple streams on a single device

RAID-S Hardware Tests

The RAID-S Symmetrix unit was also tested with `iorate` and `dd` to ensure adequate sequential bandwidth. In our tests, both mirroring and RAID-S performed very well, with few differences noted between them. At higher Symmetrix utilization levels, RAID-1 will outperform RAID-S.

In our tests, if system write limits were not reached, sequential write throughput to mirrored metavolumes was approximately 10% higher than write performance on RAID-S metavolumes. The RAID-S configuration reached write limits in some tests where the mirrored configuration did not.

EMC strongly recommends mirrored configurations for write-intensive workloads.

Observations

Fibre Channel links will show the greatest throughput if large I/O transfer sizes are used. We used large block sizes for our tests. An I/O size of 256 KB seemed to give good throughput and would be supported by most host systems. The difference in sequential read throughput with various block sizes is shown in Figure 5.

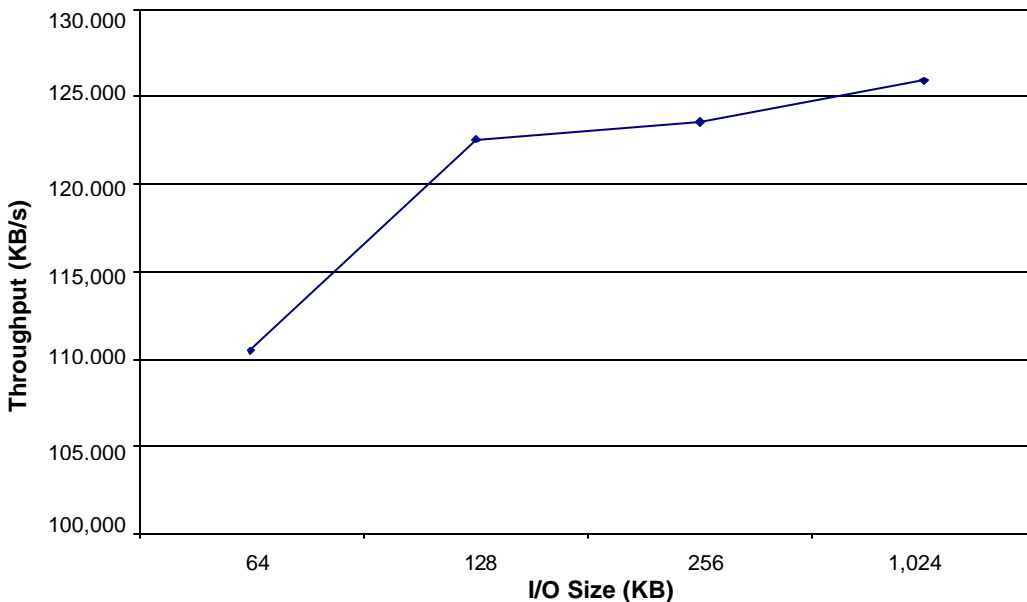


Figure 5. Sequential Read Throughput with Various Block Sizes (RAID-1)

It's important to make metavolume operations as parallel as possible. You will obtain dramatically more throughput by having multiple processes working against the metavolume. For instance, if a single thread of execution was run against a single disk, a maximum throughput of approximately 38 MB/sec was observed. When six separate threads were run against a single metavolume, we achieved approximately 132 MB/sec. With 36 threads running against a single metavolume, we achieved approximately 137 MB/sec. **This is the limit of the Fibre Channel adapters.**

EMC PowerPath was critical to the tests. The intelligent load balancing provided by PowerPath provided an advantage to both host I/O channels, even when accessing a single metavolume. Without PowerPath, throughput would have been reduced by 50%.

EMC's default metavolume stripe size (2 cylinders, or 960 KB) provided the highest throughput in our tests.

To obtain maximum throughput, it's important to ensure that each member of the metavolume is on a different disk. If multiple members of a metavolume reside on the same physical disk (or spindle), throughput will be compromised. While the EMC configuration procedures normally ensure that all members are on different physical disks, we did test several misconfigured metavolumes. These misconfigured volumes had multiple members on the same physical disk. Sequential read throughput on these volumes was approximately 55% throughput of properly configured volumes (73 MB/sec vs. 133 MB/sec).

Hardware Test Summary

In our tests, RAID-1 and RAID-S performed very similarly. The host configuration we used for the base-level hardware tests did not reach Symmetrix performance limits. EMC recommends that customers use metavolumes to simplify system management and to balance workload across multiple physical disks. The standard 2-cylinder (960 KB) stripe depth used by EMC provides optimal throughput for metavolume configurations.

Application Testing

This section describes the application testing process and the test results of building a multiple terabyte IQ-Multiplex 12.4.2 database on an EMC Symmetrix system.

Descriptions of Tables

Most tests were run against the following three tables:

- Table *CDR* consisted of 121 columns and 113 LF/HG indexes.
- Table *STR* consisted of 2 VARCHAR (255) and 1 CHAR (255) columns; all columns have FP indexes.
- Table *Dew* consisted of one date column and one integer column; the date column has an HG index, and the integer column has an HNG index.

The *CDR* table had the largest variety of columns with a mixture of indexes. One billion rows were added to the table in 100-million row increments. All data loads completed within a few seconds of each other.

The *STR* table filled up space and rows rather quickly (roughly 150 million rows/hour, adding about 85 Gbytes/hour).

Data loading into the *Dew* table was a bit faster, due to the two HG indexes. The data load rate was approximately 225 million rows/hour when loading 500-million row chunks, but it did not fill space as quickly.

Both databases and all tables were loaded using 256 KB pages/16 KB blocks.

Tables 7 and 8, respectively, provide information on the size of the total database and the size of the three tables. Table 9 summarizes the characteristics of the three tables.

Table 7. Total Database Size

| On-disk Size (compressed) (KB) | Uncompressed Size (in 256 KB pages) (logical size) | Compressed Pages | Number of blocks |
|--|---|---------------------|---------------------|
| 2,826,384,912 | 35,442,073 | 34,013,265 | 176,649,057 |
| Logical size = 35,442,073*256 = 8.8 TB. Compressed size = 2.8 TB | | | |

Table 8. Table Sizes

| Table | Compressed Size (KB) | Uncompressed Size (in 256 KB pages) (logical size) | Compressed Pages | NBlocks |
|--------------|-----------------------------|---|-------------------------|----------------|
| CDR | 275,080,330 | 1,837,640 | 1,470,112 | 17,192,521 |
| STR | 1,912,153,488 | 29,874,252 | 29,873,202 | 119,509,593 |
| Dew | 605,152,400 | 3,503,057 | 2,524,645 | 37,822,025 |

Table 9. Table Characteristics

| Table Name | Number of Rows (in Billions) | Number of Columns | Table Width (bytes) | Non-FP Indexes |
|-------------------|-------------------------------------|--------------------------|----------------------------|-----------------------|
| CDR | 1 | 121 | ~269 | 113 HG/LF |
| STR | 10.2 | 3 | 768 | none |
| Dew | 30 | 2 | 8 | 2 HG, 2 HNG |

Data I/O Rates

We used two mechanisms to test the available data I/O (or throughput) rates. We used an I/O exerciser tool built into IQ-Multiplex for use by engineering, and a query designed to maximize data throughput and minimize computation. The results are shown in Table 10.

Table 10. IOX (I/O exerciser) Results (128 KB reads)

| Disk Configuration | MB Per Second | I/O Per Second |
|--------------------|---------------|----------------|
| RAID-S | 176.7 | 1,348 |
| Mirrored | 175.9 | 1,342 |

128 KB I/O read size was chosen as similar to the average physical read for an IQ-Multiplex database that has been configured for a 256 KB page size.

During the load phase, a typical 256 KB IQ-Multiplex page gets compressed to 50% of its in-memory size prior to a disk write. This means that the I/O subsystem would be seeing 128 KB I/O requests.

Read activity dominates IQ-Multiplex activity and will typically result in a cache miss (rather than being cached within Symmetrix). This is why we focused on read activity and read performance.

The limiting factor during the IOX tests did not appear to be the CPUs, IQ-Multiplex, or the Symmetrix units. Instead, we were limited by having only two HBAs to the Symmetrix (four in total, but two were for PowerPath failover and load balancing).

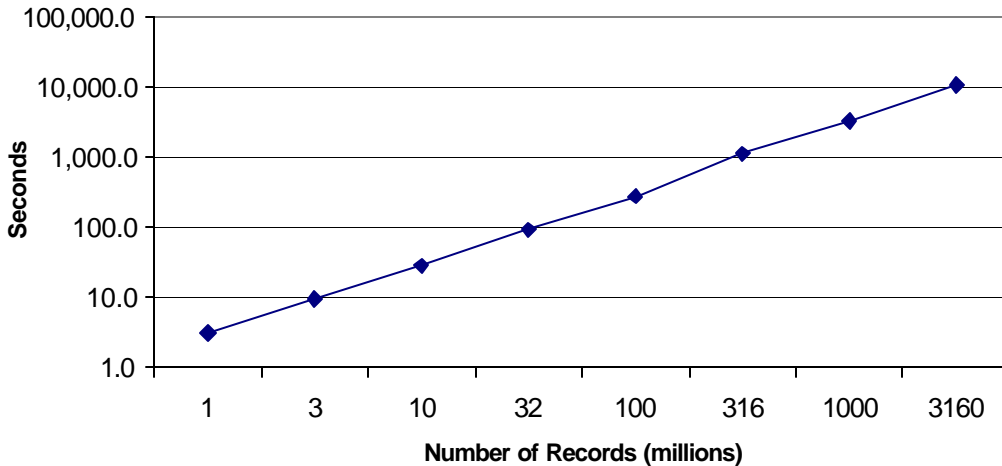
The query test was run with a four-way union of a query that selected every 25,000th cell from several columns of the CDR table. It saturated the Sun CPUs with approximately 86% user time and 14% system time. Both RAID-S and mirrored disk arrays generated about 149 MB/sec of I/O throughput.

Sorts

We ran two sort queries, and both queries sorted on a single column. The first test sorted on an integer column, while the second sorted four columns using one sort key (for a total of 1 integer + 3 CHAR(16) columns). A sort has the theoretical mathematical cost of $O(N\log N)$ or “order of N times the log of N”. We achieved test results remarkably close to this goal, which is good – it means that our sort does not get worse than it should as the size of the sort increases. Table 11 presents the 1-column and 2-column sort times by row quantity, and Figure 6 graphically presents the results of a single-column sort.

Table 11. Data Sorts – Times in Seconds

| Number of Rows (Millions) | 1-Column Sort Time (seconds) | 4-Column Sort Time (seconds) |
|---------------------------|------------------------------|------------------------------|
| 1.00 | 3.1 | |
| 3.16 | 9.5 | 20.1 |
| 10.00 | 28.6 | 63.3 |
| 31.60 | 92.4 | 223.7 |
| 100.00 | 280.2 | 825.3 |
| 316.00 | 1,124.9 | 2,950.6 |
| 1,000.00 | 3,362.6 | 9,868.8 |
| 3,160.00 | 10,565.4 | |

**Figure 6. Results of a Single-Column Sort**

For the sort, I/O rates were low: – 2.1 MB/sec reads and + 2.1 MB/sec writes for 1 column, and 12.8 MB/sec reads from main and 12.8 MB/sec writes to temp with a 4-column sort. Low I/O rates are a very good indicator of both the IQ-Multiplex data compression and of the sorting algorithm. Our implementation resulted in saturating the CPUs without overloading the I/O subsystem.

Index Performance

We tested the performance of range predicates against the date column of the *Dew* table to feed back into the optimizer. The data type is integer for all columns tested.

Our tests with a 30-billion row table (*Dew*) shows that the IQ-Multiplex engine is capable of handling very large tables.

Table 12. Results of Testing Very Large Tables (2.0-billion and 3.3-billion rows)

| Index Type | Percent of Rows selected by the Query (%) | Table with 2.0 Billion Rows (seconds) | Table with 3.3 Billion Rows (seconds) | Approximate I/O Rate and CPU Usage |
|-----------------|---|---------------------------------------|---------------------------------------|------------------------------------|
| Fast Projection | 1 | 239 | 392 | 15 MB/sec |
| | 5 | 258 | 421 | About 2 CPUs used |
| | 10 | 262 | 429 | |
| High_Non_Group | 1 | 175 | 288 | ~ 4 MB/sec |
| | 5 | 187 | 334 | Less than 1 CPU used |
| | 10 | 203 | 360 | |
| | 30 | | 467 | |
| | 50 | 315 | 543 | |
| | 80 | 319 | 548 | |
| High_Group | 1.0 | 115 | 231 | ~1.8 MB/sec |
| | 1.1 | 132 | | 1 CPU used |
| | 2.0 | 252 | 459 | |
| | 2.2 | 306 | | |
| | 3.0 | 400 | 739 | |
| | 5.37 | 818 | 1,505 | |

Application Test Summary

The host adapters are the I/O pipes, pumping data from the server (Sun Solaris in this case) to the Symmetrix unit. The more pipes you have, the more I/O you can push. The best (or most optimal) configuration would have included hosts that could accommodate at least four HBAs each; this would allow PowerPath to provide redundancy and balance the workload across all channels.

The three goals of this project were the following:

1. Validate the configuration of a VLDB IQ-Multiplex on EMC Symmetrix.
2. Validate VLDB IQ-Multiplex aspects such as 30+ billion rows in a single table
3. Perform (limited) I/O characterization of IQ-Multiplex to determine the proper sizing of the shared disk array.

Goal 1 was achieved completely. Almost 10 TB of EMC storage has been configured into IQ-Multiplex using different RAID architectures in combination with Fibre, Connectrix, and Volume Logix. A recommended practice guideline for configuring EMC Symmetrix, Sun systems, and IQ-Multiplex was developed.

Goal 2 proved the ability of IQ-Multiplex to load and query a multiterabyte database. This project effectively used a variety of EMC hardware and software products to provide a highly scalable, good performing, and reliable solution that can satisfy any/all data warehousing needs and requirements.

In regard to Goal 3, a number of I/O tests were performed to verify Sybase sizing rules for IQ-Multiplex. While the standard IQ-Multiplex sizing rule of 10 MB/sec per CPU has been verified and even proven conservative, tests indicated that some environments may show different I/O behavior, mostly needing less than 10 MB/sec per CPU, while some may need more I/P bandwidth per CPU. Tests also validated that EMC shared-disk arrays will allow scalability of IQ-Multiplex well into multi-CPU range.

Additional Documentation

To obtain database related white papers, contact your local EMC representative. White papers can be found on the following internal web page: <http://esgw3/ewiz>

For more information on Sybase IQ-Multiplex and other Sybase Business Intelligence solutions, visit the Sybase Business Intelligence web page at <http://www.sybase.com/>.

Many thanks to the Sybase technical staff who worked on this project and contributed to this paper.