



HP reference configuration for SQL Server 2008 Fast Track Data Warehouse: ProLiant DL585 and StorageWorks MSA2000fc

Table of contents

Executive summary.....	2
Fast Track reference architecture.....	2
Important points and caveats.....	2
Solution criteria	4
Storage options	5
Bill of materials.....	7
For more information.....	9

Executive summary

Microsoft® SQL Server Fast Track Data Warehouse for Hewlett-Packard (HP) servers, storage and networking products provides a prescriptive approach for balancing server, storage, network and software configurations for architecting Microsoft SQL Server 2008 data warehouse solutions. The reference architectures provide server and storage guidance for various data warehouse workloads—giving you the most efficient hardware for your solution, saving you time and cost in choosing the right technology, and giving you peace of mind that the right platform and architecture is in place.

Target audience: The target audience for this document consists of IT planners, architects, DBAs, CIOs, CTOs, and business intelligence (BI) users with an interest in options for their BI applications and in the factors that affect those options.

This white paper describes testing performed by HP in September 2009.

Fast Track reference architecture


This document is a reference configuration companion document to the Microsoft SQL Server Technical Article, “Implementing a SQL Server Fast Track Data Warehouse”, which describes a repeatable architectural approach for implementing a scalable model for a symmetric multiprocessor (SMP)-based Microsoft SQL Server 2008 data warehouse. The end result of the process described in this companion guide represents a recommended minimal SQL Server 2008 configuration, inclusive of all the software and hardware, required to achieve and maintain a baseline level of “out of box” scalable performance when deploying SQL Server 2008 data warehousing (SSDW) sequential data access workload scenarios versus traditional random I/O methods.

This document provides specific details about the configuration and bill of materials for one such reference architecture, based on the HP ProLiant DL585 G6 and the HP StorageWorks 2000fc Modular Smart Array G2 (MSA2000fc). This configuration is targeted at the middle tier of a data warehousing environment, targeting moderately complex queries and concurrency. It is optimized at 12 TB of compressed storage and is designed to scale to 24 TB of user data capacity.

Important points and caveats

- The configuration described here and the approach detailed in the reference architecture guide is exclusively designed for, and is only applicable to, sequential data workloads. Use of this approach on other workload types is not appropriate and may yield configurations that are inefficient.
- **ALL** recommendations and best practices defined in the reference architecture guide must be implemented in their entirety in order to preserve and maintain the sequential order of the data and sequential I/O against the data.
- The recommendations detailed in this paper have been reached through recent lab-based testing against benchmark databases, performed by HP on HP hardware. Microsoft and HP plan to update the reference architecture guide and this companion document with additional practices to help maintain the optimal data layout under more diverse conditions, and any applicable configuration changes, as the use of these architectures becomes more widespread.

Figure 1. Fast Track 12TB reference architecture with the ProLiant DL585 G6 and (6) StorageWorks MSA2312fc



A complete SQL Server DBMS (Database Management System) configuration, or “stack,” is a collection of all the components that are configured to work together to support the database application. This includes the physical server hardware (with its BIOS settings and appropriate firmware releases), memory, CPU (number, type, clock and bus speed, cache and core count), operating system settings, the storage arrays and interconnects, disk (capacity, form factor and spindle speeds), database, DBMS settings and configuration, and even table types, indexing strategy, and physical data layout.

The primary goal of Fast Track, which is also a common goal when designing most data center infrastructures, is a balanced configuration where all components can be utilized to their maximum capability. Architecting and maintaining a balance prevents over subscribing certain components within the stack to a point where the expected performance is not realized; understanding the performance limits of your configuration can help prevent wasted cost for components that will never realize their potential due to other constraints within the stack.

Solution criteria

The reference configurations are built on three HP ProLiant server platforms, each targeting a different tier of a SQL Server Data Warehousing solution. The Microsoft and HP Fast Track architectures gain greater throughput and scale by using the following approach:

- Targeting query workloads patterned for large sequential data sets rather than small random data transactions
- Optimizing rapid data reads and query aggregations

All three configurations leverage the HP StorageWorks MSA2000fc array, which allows for dual reads when drives are mirrored. For sequential data reads from data warehouse queries, this capability enables tremendous throughput per storage volume—up to 240 MB/s. The Fast Track approach, and supporting storage array architecture, is optimized for sequential reads. To support a non-optimized, random I/O data warehousing workload, up to 2 to 3 times the number of drives would be required to achieve the same throughput.

The ProLiant DL385 G6 configuration is ideal for data marts with a small departmental footprint of query activity with a volume optimized at 6 TB of compressed data and scalable to 12 TB of user data capacity. The ProLiant DL585 G6 configuration serves the middle tier of data warehouse activity, targeting moderately complex queries and concurrency. It is optimized at 12 TB of compressed storage and scalable to 24 TB of user data capacity. The ProLiant DL785 G6 configuration, with eight 6-core processors, is the top scale-up model in the ProLiant server line, capable of handling higher concurrency and query complexity. It is optimized at 24 TB of compressed storage and scalable to 48 TB of user data capacity. Table 1 summarizes these configurations.

Table 1. Fast Track reference configurations with HP ProLiant servers and StorageWorks arrays

Server	CPU	Total Cores	SAN	Drive Count (tested capacity)	Tested Capacity	Max Capacity
DL385 G6	(2) AMD Opteron™ Model 2435	12	(3) HP MSA2000fc	(24) 300GB 15k SAS	6TB	12TB
DL585 G6	(4) AMD Opteron Model 8435	24	(6) HP MSA2000fc	(48) 300GB 15k SAS	12TB	24TB
DL785 G6	(8) AMD Opteron Model 8439 SE	48	(12) HP MSA2000fc	(96) 300GB 15k SAS	24TB	48TB

The focus of this reference configuration document is the data mart-focused configuration, based on the ProLiant DL585 G6. Table 2 outlines the configuration details for this platform.

Table 2. ProLiant DL585 G6 configuration details

Model	ProLiant DL585 G6 8435 4P
CPU	(4) Six-Core AMD Opteron Processor Model 8435 (2.6 GHz, 6MB L3 cache, 75W)
Number Cores	24
PCI-E Slots	(3) x8, (4) x4
PCI-X Slots	(2) 100MHz
Drives	(2) HP 72GB 3G SAS 15K SFF DP ENT HDD (6) HP 300GB 3G SAS 10K SFF DP ENT HDD
Storage Controller	HP Smart Array P410/512 BBWC with Heat Sink
Host Bus Adapters	(4) HP 82E 8Gb Dual Port PCI-E FC HBA
Network Adapters	Two HP NC371i Multifunction Gigabit Server Adapters with TCP/IP Offload Engine, including support for Accelerated iSCSI through an optional Licensing Kit (four ports total)
RAM	96 GB PC2-6400 expandable to 256GB

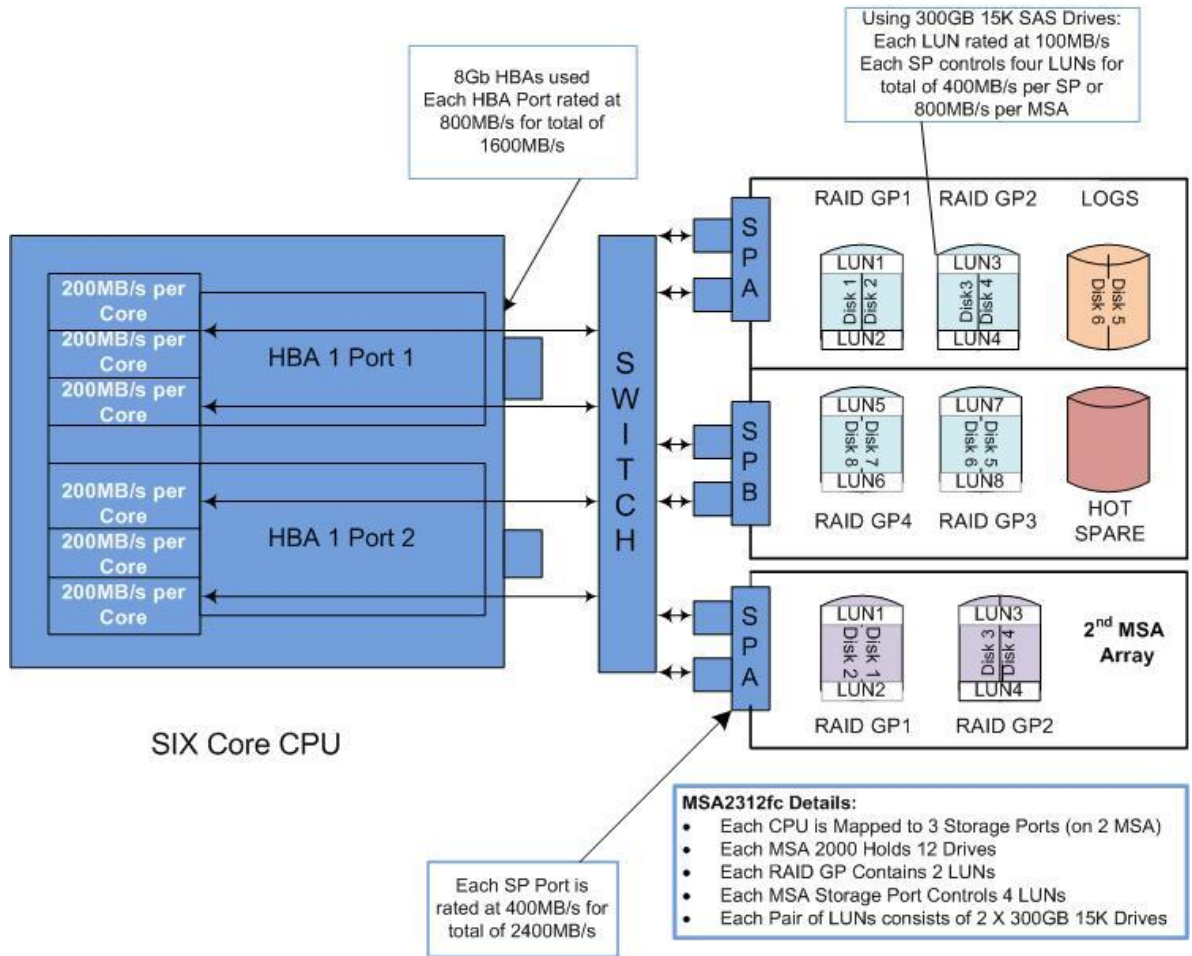
Note that the recommended 96GB is the minimum RAM for this configuration. Generally, for SQL Server Data Warehousing environments, as workload demands grow, increasing RAM provides performance benefits. If the workload consists of a large number of small-to-medium sized queries hitting the same data (for example, last week's sales), performance and throughput can be increased by caching results in memory.

Storage options

The Fast Track reference configuration is based on a model of a core-balanced architecture which allows for a configuration that is both initially balanced and easily scalable in sized "building blocks," each providing predictable performance. These building blocks are sized for an optimum number of storage components to drive a certain number of CPU cores. The diagram below represents the building blocks for each six-core processor of the HP ProLiant DL585 G6 using HP MSA2000fc storage arrays required to achieve core-balanced sequential data I/O (regardless of capacity).

Figure 2. Six-core processor balanced architecture

**SQL Server 2008 Minimum Server Configuration
SMP Core-Balanced Architecture using Dual Read on HP MSA2000fc**



SP=Storage Port, GP=Group

More information regarding the Fast Track Data Warehouse reference architecture approach, including the means of calculating building blocks for other HP server and storage combinations, is included in the Microsoft SQL Server Technical Article, "Implementing a SQL Server Fast Track Data Warehouse".

It is also possible to increase storage capacity without adding more CPU cores by using larger drives or more drives per LUN; however increasing the capacity beyond the optimized configuration will not increase performance. Rather, only capacity will be increased. Important to note is that adding more data but still querying the same sized ranges per query will not decrease performance. Table 3 summarizes the configuration specifics and capabilities of the fibre channel version of the MSA2000, the MSA2312fc.

Table 3. MSA2000fc storage configuration details

Model	(6) MSA2312fc Fibre Channel Array with dual controllers
Drives	(66) 300GB 15K LFF SAS disks
Drive layout	(48) drives for user data (8 per MSA), configured as (24) 2-disk RAID1 arrays (12) drives for log data (2 drives each MSA), configured as (6) 2-disk RAID1 arrays (6) drives for hot spares (1 per MSA)
Raw space (user data)	6.67TB
Est. User data capacity	4.8 TB (no compression) 12 TB (with compression) Notes: 1) Compression is assumed to be 2.5:1 2) Drive sizes can be changed to increase or decrease storage capacity 3) An additional MSA2000 drive enclosure can be added to the MSA2312fc to increase capacity
Est. Log Capacity	1,674 GB

Bill of materials

Table 4 includes the bill of materials for the server, storage, switching and rack components of the reference configuration.

Table 4. Bill of materials

Qty	Part Number	Description
Server Configuration		
1	574409-B21	HP ProLiant DL585 G6 Rack CTO Chassis
1	462864-B21	HP SA P410/512 BBWC
1	572373-L21	AMD 8435 6 core 2P option kit
1	572373-B21	AMD 8435 6 core 585 G6 processor option kit
2	418371-B21	HP 72GB 3G SAS 15K SFF DP ENT HDD
6	492620-B21	HP 300GB 3G SAS 10K SFF DP ENT HDD
4	AJ763A	HP 82E 8Gb Dual Port PCI-E FC HBA
12	497767-B21	8 GB REG PC2-6400 2 x 4 GB Kit
Storage Configuration		
6	AJ795A	HP 2312fc DC Enh Modular Smart Array
66	AJ736A	HP MSA2 300GB 15k rpm 3.5 SAS HDD

Qty	Part Number	Description
Switch Configuration		
1	AM869A	HP 8/40 40-Ports / 24-Active Enabled SAN Switch
1	T5519A	HP 8/40 SAN Switch 8-Port Upgrade Activation License
32	221692-B21	2m LC-LC Cable
2	263474-B22	HP Cat 5e Cables – 8 Pack
24	AJ715A	HP 4Gb Shortwave B-series FC SFP+ 1 Pack
8	AJ716A	HP 8Gb Shortwave B-series FC SFP+ 1 Pack
1	J9279A	HP ProCurve Switch 2510G-24
Rack Configuration		
1	AF002A	HP Universal Rack 10642 G2 Shock Rack
1	AF002A#001	Factory Express Base Racking
1	AF062A	HP 10K G2 600W Stabilizer Kit
1	AF054A	HP 10642 G2 Sidepanel Kit
1	AF054A#0D1	Factory integrated
2	252663-B24	HP 16A High Voltage Modular PDU
1	252663-B24#0D1	Factory integrated
2	AF593A	HP 3.6m C19 Nema L6-20P NA/JP Pwr Crd
1	AF593A#0D1	Factory integrated
1	120672-B21	HP 9000 Series Ballast Option Kit
1	120672-B21#0D1	Factory integrated

For more information

Implementing a SQL Server Fast Track Data Warehouse, <http://msdn.microsoft.com/en-us/library/dd459178.aspx>

HP ActiveAnswers, www.hp.com/solutions/activeanswers

SQL Server solutions on ActiveAnswers, www.hp.com/solutions/activeanswers/microsoft/sql

HP and Microsoft SQL Server Business Intelligence solutions, www.hp.com/solutions/microsoft/sqlbi

HP and Microsoft SQL Server Business Intelligence Configurations and SQL Server Fast Track Data Warehouse reference architectures, www.hp.com/solutions/microsoft/sqlbiconfigs

HP and Microsoft SQL Server solutions, www.hp.com/solutions/microsoft/sql

HP and Microsoft, www.hp.com/go/microsoft

HP ProLiant DL585 G6 server, www.hp.com/servers/dl585-g6

HP StorageWorks MSA2000, www.hp.com/go/msa

HP ProCurve Networking, www.procurve.com

How to buy, www.hp.com/buy

To help us improve our documents, please provide feedback at http://h20219.www2.hp.com/ActiveAnswers/us/en/solutions/technical_tools_feedback.html.

Technology for better business outcomes

© Copyright 2009 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft is a U.S. registered trademark of Microsoft Corporation. AMD Opteron is a trademark of Advanced Micro Devices, Inc.

4AA2-4631ENW, Revision 3, October 2009



Get connected

www.hp.com/go/getconnected

Current HP drivers, support & security alerts
delivered directly to your desktop

