Microsoft SQL Server 2012 Fast Track Reference Configuration Using PowerEdge R720 and EqualLogic PS6110XV Arrays

This whitepaper describes Dell | Microsoft SQL Server Fast Track reference architecture configurations and performance results

Jisha J

Dell Database Solutions Engineering

Chris Boire

Dell EqualLogic Engineering



This document is for informational purposes only and may contain typographical errors and technical inaccuracies. The content is provided as is, without express or implied warranties of any kind.

© 2012 Dell Inc. All rights reserved. Dell and its affiliates cannot be responsible for errors or omissions in typography or photography. Dell, the Dell logo, PowerEdge, and EqualLogic are trademarks of Dell Inc. Intel and Xeon are registered trademarks of Intel Corporation in the U.S. and/or other countries. Microsoft, Windows Server, and SQL Server are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

December 2012 | Rev 1.0

Contents

Executive summary	.4
Introduction to FTDW reference architectures	.4
What's new with FTDW reference architectures?	.4
Reference architectures	.5
Single-server reference architecture	. 5
Highly-available reference architecture	.7
Hardware details	. 8
Configuration tuning details	10
Additional storage requirements for the high-availability (HA) configuration	15
Performance details	16
Summary	17

Tables

Figure 1: Single-server reference architecture	.6
Figure 2: Highly-available reference architecture	.7
Figure 3: Cabling diagram for single-server configuration	10
Figure 4: MPIO Settings for network adapter ports	11
Figure 5: Enabling jumbo packets on network adapter ports	12
Figure 6: Increasing Transmit Buffers on network adapter ports	13
Figure 7: Force10 S4810 Switch Settings	14
Figure 8: Storage LUN details	15

Figures

Figure 1: Single-server reference architecture	. 6
Figure 2: Highly-available reference architecture	.7
Figure 3: Cabling diagram for single-server configuration	10
Figure 4: MPIO Settings for network adapter ports	11
Figure 5: Enabling jumbo packets on network adapter ports	12
Figure 6: Increasing Transmit Buffers on network adapter ports	13
Figure 7: Force10 S4810 Switch Settings	14
Figure 8: Storage LUN details	15

Executive summary

Dell and Microsoft cooperate to provide guidelines and design principles that assist customers in designing and implementing balanced configurations for Microsoft® SQL Server® data warehouse workloads to achieve "out-of-the-box" scalable performance. These database reference architectures enable each of the components in a database stack to provide optimal throughput to match the capabilities of the specific setup.

This white paper describes architectural design principles and guidelines to achieve optimally balanced data warehouse solutions using Dell[™] PowerEdge[™] R720 servers and Dell EqualLogic[™] PS6110XV iSCSI SAN storage with Microsoft SQL Server 2012 in single-server and highly-available configurations. The paper also provides performance data for these reference configurations.

Introduction to FTDW reference architectures

The objective of Fast Track Data Warehouse (FTDW) reference architectures is to provide tested and validated configurations for optimized data warehouse solutions. These reference architectures provide balanced server, memory, network, and storage configurations that help ensure system capability and throughput is maximally utilized for data warehouse and business Intelligence (BI) systems.

Benefits of following these best practices and guidelines include:

- Accelerate data warehouse projects with pre-tested hardware and SQL Server configurations.
- Reduce hardware and maintenance costs by purchasing a balanced hardware solution and optimizing it for a data warehouse workload.
- Reduce planning and setup costs by leveraging certified reference architecture configurations.
- Predict performance by implementing an optimized configuration and taking advantage of tuning recommendations.

Visit <u>msdn.microsoft.com/en-us/library/dd459146%28v=sql.100%29.aspx</u> for more information on Fast Track Data Warehouse architectures.

What's new with Dell | Microsoft FTDW reference architectures?

Dell and Microsoft recently updated FTDW reference architecture offerings with the latest technology advancements in database, server, and storage technologies. Reference architectures and configurations are jointly engineered by Dell and Microsoft—hardware and software optimizations are tested by Dell, and performance results are cross-checked by Microsoft. This approach offers faster time-to-value for customers using integrated, balanced, and verified architectures.

One of the many new features included in Microsoft SQL Server 2012 is Column Store Index. This feature enables storing data in columnar fashion, in contrast to the traditional row-based approach, to provide better database compression rates. This approach is especially beneficial for data warehouses, due to the huge amount of data handled. Column Store Indexes also benefit common data warehousing queries such as filtering, aggregating, grouping, and star-join queries (see <u>msdn.microsoft.com/en-us/library/gg492088.aspx</u> for details).

Dell PowerEdge 12th generation servers feature enhanced onboard memory, storage, and processor speeds, and include other advanced features that improve data warehouse performance. Additionally, they include the latest Intel[®] Xeon[®] E5 series processors, larger memory capacities, higher memory speed, and third-generation PCI Express slots to ensure faster database throughput.

EqualLogic storage arrays offer enterprise-class performance and reliability, intelligent automation, and seamless virtualization of storage with simplified storage management. The EqualLogic PS Series offers exceptional performance for both sequential and transactional applications, with linear scalability as arrays are added. It delivers a modular and cost-effective solution that can be deployed in appropriate increments for small and medium businesses.

Reference architectures

This paper proposes two different FTDW reference architectures comprised of SQL Server 2012, PowerEdge R720 servers, and EqualLogic PS6110XV storage arrays. Table 1 lists the Proposed Reference architectures, along with the assigned Solution IDs.

Reference Architecture	Solution ID
Single Server Reference Architecture	2813604
Highly Available Reference Architecture	2813607

Table 1: Recommended reference architectures with Solution IDs

Single-server reference architecture

The single server reference architecture comprises of the minimal and optimal configuration required to setup the SQL Server fasttrack reference architecture. The specific configuration along with the software and hardware tuning implementations is certified as Dell Microsoft SQL Server fasttrack configuration.

Figure 1 depicts major elements of the single-server reference architecture. Table 2 lists details of the configuration.

Figure 1: Single-server reference architecture



Table 2: Single-server reference architecture details

Hardware Components	Details
Server	One PowerEdge R720 server
Processors	Two Intel® Xeon® Sandy Bridge CPUs E5-2640 (2.5Ghz)
Total Cores per Socket	6
Total Logical Processors (HT enabled)	24
Total Installed Memory	128GB @ 1333Mhz
Network Adapters	 SAN Connectivity: Two Intel Ethernet 10G 2P X520 Adapters Public Connectivity: Minimum of one network adapter (1 Gbps or 10Gbps based on the requirement); it is recommended to have more than one network adapter with load balancing configured
Multipathing Software	EqualLogic MPIO DSM
Multipathing Policy	Least Queue Depth (Default)
External Storage	Three EqualLogic PS6110XV arrays

Disks	Seventy-two 146G, 6 Gbps 15k SAS drives (Including six hotspare drives)
Network Switch	One Dell Force10® S4810 switch
Operating System	Microsoft Windows Server® 2008 R2 SP1 Enterprise Edition
Database Software	SQL Server 2012 Enterprise Edition

Highly-available reference architecture

For achieving high database availability, it is recommended to use Microsoft Database Clustering. Using Microsoft clustering services, one database server is configured as the primary (active) server, and the second server is configured as secondary (passive). The secondary server should have exactly the same configuration as the primary server. Since the database is active on a single server at any point in time, the performance of the database on the primary server (active) is comparable to the single server configuration (discussed earlier).

Figure 2 depicts the highly available fast track reference architecture (with the major components called out), and Table 3 lists the configuration details.

Figure 2: Highly-available reference architecture



Table 3: Highly-available reference architecture details

Hardware Components	Details
Server	Two PowerEdge R720 servers
Processors	Two Intel Xeon Sandy Bridge CPUs (E5-2640 @ 2.5Ghz) per PowerEdge server
Total Cores per Socket	6
Total Logical Processors (HT Enabled)	24 per server

Total Installed Memory	128GB @ 1333Mhz per server	
Network Adapters	 SAN Connectivity: Two Intel Ethernet 10G 2P X520 Adapter per PowerEdge server Public Connectivity: Minimum of one network adapter (1Gbps or 10Gbps, based on the requirement) per server; it is recommended to have more than one network adapter with load balancing configured per server Private Network (Cluster): Minimum of one network adapter with 1Gbps (or more) connectivity per server 	
Multipathing Software	EqualLogic MPIO DSM	
Multipathing Policy Least Queue Depth(Default)		
External Storage	Three Dell EqualLogic PS6110XV arrays	
Disks	Seventy-two 146G, 6Gbps 15k SAS drives (Including six hotspare drives)	
Network Switch	Two Dell Force10 S4810 switches	
Operating System	Windows Server 2008 R2 SP1 Enterprise Edition	
Clustering Software	Microsoft Windows Clustering	
Cluster Quorum Configuration	Node and Disk Majority (Recommended)	
Database Software	SQL Server 2012 Enterprise Edition	

Hardware details

This section provides details about the hardware used in the reference architectures.

Dell PowerEdge R720 server

The PowerEdge R720 is a 2-socket 2U rack server that offers simplified management, purposeful design, and energy efficiency, and supports Intel E5 Series Sandy-Bridge processors and ECC DDR3 RDIMMs with a maximum capacity of 768GB Memory. PowerEdge R720 supports five x8 PCIe Gen3 and two x16 PCIe Gen3 slots.

Visit <u>dell.com/us/enterprise/p/poweredge-r720/pd</u> for more information on PowerEdge R720 Servers.

Intel Ethernet 10G 2P X520 server adapter

The proposed reference architectures use two dual-port Intel Ethernet 10G 2P X520 server adapter cards. This dual-port adapter has RJ-45 copper interfaces, and is available both in low-profile and full-height interfaces.

Visit <u>intel.com/content/www/us/en/network-adapters/10-gigabit-network-adapters/ethernet-</u> <u>10gigabit-adapters.html</u> for more information on Intel Ethernet 10G 2P X520 server adapters.

Dell Force10 S4810 Ethernet switch

The Dell Force10 S4810 switch is a 1U rack-mountable 10/40Gb Ethernet switch with 48 ports. It features a 1.28Tbps (full-duplex) non-blocking switching fabric designed to deliver line-rate performance under full load with low application latency.

Visit <u>dell.com/us/business/p/force-10-s-series/pd?c=us&s=bsd&~cd=anave</u> for more information on Dell Force10 S-series switches.

EquallLogic PS6110XV iSCSI storage arrays

EqualLogic PS6110XV storage arrays support 2.5-inch, 6Gbps SAS disks that enable high backend performance. Each PS6110 array supports up to two controllers in active-passive failover, and each controller has a single 10GbE host-side port. Each controller module contains 2GB of cache that mirrors the other controller's cache for high availability and is protected by a battery-powered cache offload mechanism. EqualLogic peer storage arrays provide virtualized storage—host volumes are carved out of storage pools consisting of one to four arrays. Automatic optimization mechanisms ensure high performance and availability by distributing volume pages within and across arrays.

Dell EqualLogic storage comes with its own Device Specific Module (DSM) software, which is fully integrated with Microsoft MPIO to help configure multipath solutions.

Figure 3 depicts how cables are connected in the configuration.



Figure 3: Cabling diagram for single-server configuration

Visit <u>dell.com/us/enterprise/p/equallogic-ps6110-series</u> for more information on Dell EqualLogic PS6110XV storage arrays.

Configuration tuning details

This section explains best practices and tuning guidelines to achieve optimal performance in the discussed reference architectures.

EqualLogic PS6110XV storage tuning

EqualLogic peer storage arrays are self-optimizing; non-default settings are not required to obtain maximum sequential throughput. Instead, storage arrays automatically detect and adapt to configuration choices made on the host and switch—for example:

- Jumbo Frames supported on all storage iSCSI host ports (adaptive)
- Single iSCSI session per volume slice per array (default)

Figure 4 shows a snapshot of the EqualLogic Remote Setup Wizard to help verify multipathing settings.

Remote Setup Wize	ard			×
	Configure MPIO settings	for this comp	outer	
	Subnets included for MPIO		Subnets excluded from MPIO	
	10.124.54.0 - 255.255.254.0 169.254.0.0 - 255.255.0.0	Exclude >>		
	Default Load Balancing Policy-	Least Queue [Depth C Round Robin	
	Max sessions per volume slice:	1	Use MPIO for snapshots	
	Max sessions per entire volume:	3 🗧 🤇	Use IPv4 O Use IPv6	
	Minimum adapter speed:	1 Gbps	•	
	< Back	Next >	Finish	

Figure 4: MPIO Settings for network adapter ports

Intel X520 10GbE NIC tuning

Jumbo packets are enabled on all NIC ports and set to a packet size of 9014 bytes. Jumbo packet settings can be changed on the **Advanced** tab of the Adapter Properties window (Figure 5).

Intel(R) Ethernet 10G 2P X520 Adapter Properties
VLANs Boot Options Driver Details General Link Speed Advanced Teaming
Advanced Adapter Settings
Profile: Standard Server
Settings: Value:
Interrupt Moderation Jumbo Packet Large Send Offload V2 (IPv4) Large Send Offload V2 (IPv6) Locally Administered Address Log Link State Event Maximum Number of BSS Processors ▼ Use Default
Jumbo Packet
Enables Jumbo Packet capability for TCP/IP packets. In situations where large packets make up the majority of traffic and additional latency can be tolerated, Jumbo Packets can reduce CPU utilization and improve wire efficiency. Jumbo Packets are larger than standard Ethernet frames, which
are approximately 1.5k in size. Image: Note: Changing this setting may cause a momentary loss of connectivity.
OK Cancel

Figure 5: Enabling jumbo packets on network adapter ports

Also on the **Advanced** tab of the Adapter Properties window, under Performance Options Properties, the number of Receive and Transmit Buffers are increased to the maximums of 4096 and 16384, respectively.

Performance Options	×		
Settings:	Value:		
Flow Control Interrupt Moderation Rate Low Latency Interrupts Receive Buffers Transmit Buffers	16,384		
	Use Default		
Transmit Buffers			
Sets the number of Transmit Buffers used by the adapter when copying data to memory. Increasing this value can enhance transmission performance, but also consumes system memory.			
You might choose to increase the number of Transmit Buffers if you notice a significant decrease in the performance of transmitted traffic. If transmission performance is not an issue, use the default setting.			
	OK Cancel		

Figure 6: Increasing Transmit Buffers on network adapter ports

Force10 S4810 Switch Tuning

It is recommended to use the following settings on all participating iSCSI ports of the configuration:

- Set all ports to 10Gb-per-second Ethernet
- Enable switchport mode
- Enable receive and transmit flow control
- Enable jumbo packets
- Set all ports as edge ports

Figure 7 shows an example of how to implement these settings from the switch console.

```
Figure 7: Force10 S4810 Switch Settings
```

```
This is the Force10 S4810 switch stack in the PerfPirate Rack

IP address: 10.126.6.30

Login: admin

Password:

This is the Force10 S4810 switch stack in the PerfPirate Rack

IP address: 10.126.6.30

PerfPirate10g-sw0@confj@ure

PerfPirate10g-sw0(conf)#interface range tengigabitethernet 0/0 - 47

PerfPirate10g-sw0(conf)if-range-te-0/0-47)#switchport

PerfPirate10g-sw0(conf-if-range-te-0/0-47)#flowcontrol rx on tx on

PerfPirate10g-sw0(conf-if-range-te-0/0-47)#mtu 9252

PerfPirate10g-sw0(conf-if-range-te-0/0-47)#spanning-tree rstp edge-port

PerfPirate10g-sw0(conf)#exit

PerfPirate10g-sw0(conf)#exit

PerfPirate10g-sw0[conf)#exit

PerfPirate10g-sw0[conf)#exit

PerfPirate10g-sw0[conf)#exit
```

Windows tuning

The allocation unit size for all the database volumes is set to 64KB.

SQL Server 2012 tuning

The following startup options are added to the SQL Server Startup options:

- -E: This parameter increases the number of contiguous extends in each file that are allocated to a database table as it grows. This option is beneficial because it improves sequential access.
- **-T1117**: This trace flag ensures the even growth of all files in a file group when auto growth is enabled. Note that the Fast Track reference guidelines recommend pre-allocating the data file space rather than depending on auto grow.
- SQL Server Maximum Memory: For SQL Server 2012, Fast Track 4.0 guidelines suggest allocating no more than 92% of total server RAM to SQL Server. If additional applications will share the server, the amount of RAM left available to the operating system should be adjusted accordingly. For this reference architecture, the maximum server memory was set at 117.76GB.

For the specific recommended fast track configuration, the resource governor (disabled by default) and the 'max degree of parallelism' settings (0 by default) are set to the default settings.

Storage system configuration

FTDW reference architecture guidelines define three primary layers of storage configuration:

- Physical disk array (RAID Groups for data and logs)
- Operating system volume assignment (LUN)
- Databases: User, System Temp, Log

EqualLogic PS6110XV arrays are fully-virtualized storage components where the data placement of the volumes is decided by the array intelligence.

In the recommended Fast Track configurations, three PS6110XV arrays are grouped into a single storage group, and all are assigned to the same default storage pool. All arrays are configured using RAID 50, and the volumes are created from the default storage pool. All volumes span all arrays, and the portion of the volume on any one array, called a volume slice, is owned by the active storage controller of that array.

- For the primary user and system temp database files, eight 360GB volumes (LUNs) are created from the storage pool.
- For the primary user and system temp database transaction logs, one 360GB volume is created. This RAID group is dedicated to host the database transaction log files.

By default, two of the disks per EqualLogic array are dedicated as hot spares. Therefore, a total of 66 disks are used for the database configuration.

Figure 8 shows a data LUN in the EqualLogic group manager interface.

父 Volume vol01		🗞 C 🔇 📀 ?
Activities	Status Access Snapshots Replication Collections Schedules Connections	
😝 Volume vol01 🛛 🔺	General Volume Information	?
Volume Modify settings Clone Set offline Set access type	Status: O online General Settings ISCSI access: restricted Volume name: vol01 Access type: read-write, not shared Reported size: 360 GB ISCSI connections: 3 3	Replication: disabled SyncRep: disabled
 <u>Delete volume</u> <u>Convert to template</u> 	Volume and Snapshot Space	•
Move volume Folder	Storage pool default	
Move to folder Access Create access policy Snapshots	Volume Space Reported size: 380 GB In-use 360 GB Thin provisioning: disabled In-use 360 GB Volume reserve: 380 GB (100%) In-use 360 GB	
Modify snapshot settings Create snapshot Restore volume Delete snapshots	Snapshot Space Snapshot reserve: 0 MB (0%) Space recovery: delete oldest snapshot Number of snapshots: 0	100%
Schedules Create schedule Replication Configure replication	Volume Distribution RAID preference: automatic perfpirate8: 119.96 GB perfpirate6: 120.16 GB perf Load balancing: enabled perfpirate8: 119.96 GB perfpirate6: 120.16 GB perf	pirate7: 119.88 GB
SyncRep Configure SyncRep		

Figure 8: Storage LUN details

As in Figure 8, each of the LUNs to be evenly distributed across all storage arrays.

It is recommended to use mount points for the storage access, rather than using drive letters. It is also important to assign the appropriate LUN volume and mount point names to the configuration to simplify troubleshooting and performance analysis. All logical volumes are mounted to the C:\FT folder.

Additional storage requirements for the high-availability (HA) configuration

The proposed HA reference configuration leverages Windows Server Failover Clustering (WSFC) to achieve high availability. The following storage-side configuration details are required to configure clustering:

- Quorum Configuration: To increase the availability of the cluster and the SQL server database, it is important to set the cluster quorum (disk witness) configuration appropriately.
 - Create a storage LUN (of at least 512MB size) out of the storage pool, and assign it to the cluster to be configured as the disk witness (quorum). It is recommended to have 2GB of quorum configured for the proposed architecture.
 - Set the quorum mode to 'Node majority with witness' (Node and Disk Majority) for the two node cluster configuration.

Visit <u>technet.microsoft.com/en-us/library/cc731739.aspx</u> for more information on configuring Quorum in a Failover Cluster.

• Add all of the Storage LUNs used in the reference configuration to the cluster resource group.

Performance details

Table 4 shows the performance data reported for the recommended reference architectures.

Metric	Value	Description
FTDW Rated Capacity	15 TB	This capacity rating is based on "up-to" capacity, but adjusted to account for Fast Track Rated I/O.
FTDW Rated I/O	2021.25 MB/s	Core performance metric for validation. This is the midpoint of physical and logical I/O.
Up-to User Data Capacity	23 TB	Maximum user data capacity. Includes the total disk capacity of all disks allocated to primary data storage, and assumes a SQL Server page compression ratio of 3.5:1. This metric is not limited by rated bandwidth.
Benchmark Scan Rate Logical	2438 MB/s	Reflects actual user query throughput, which includes reads from RAM/Buffer cache.
Benchmark Scan Rate Physical	1604.5 MB/s	Reflects physical I/O read from disk during benchmark.
FTDW Peak I/O	2804 MB/s	Maximum observed I/O rate.
FTDW Rated CSI	4637 MB/s	Represents potential throughput using Columnstore Index.

Table 4: Performance metrics

Summary

Dell | Microsoft FTDW solutions enable customers to deploy efficient and reliable data warehouse systems using validated architectures that ensure optimally designed BI solutions. These best practices and recommendations enable customers to achieve enhanced Return on Investment (ROI) in a balanced data warehouse environment, with better performance than traditional data warehouse systems.

Dell | Microsoft Fast Track architectures provide the following benefits:

- Deliver a tested and validated configuration with proven methodology and performance behavior.
- Achieve a balanced and optimized system at all the levels of the stack by following the best practices of hardware and software components.
- Avoid over-provisioning of hardware resources.
- Offer high availability at all levels of setup (host, switches, and storage).
- Help customers avoid the pitfalls of an improperly designed and configured system.
- Reduce future support costs by limiting solution re-architect efforts because of scalability challenges.