Dell EqualLogic Best Practices Series

Dell EqualLogic PS Series Reference Architecture for Cisco Catalyst 3750X

Two-Switch SAN Reference

Storage Infrastructure and Solutions Engineering
Dell Product Group
April 2012

Table of Contents

Acknowledgements

This whitepaper was produced by the PG Storage Infrastructure and Solutions team between January 2011 and April 2011 at the Dell Labs facility in Round Rock, Texas.

The team that created this whitepaper:

Nirav Shah, Tony Ansley, and **Margaret Boeneke**

Feedback

We encourage readers of this publication to provide feedback on the quality and usefulness of this information. You can submit feedback to this email address:

SISfeedback@Dell.com

1 Introduction

The Cisco® Catalyst 3750X-48 is a 48 port, 1 Gigabit (Gb) Ethernet workgroup switch with 10G LAG ports as well as dedicated stacking interfaces.

Testing has been performed that proves the value of the Catalyst 3750X switch when used to build a small branch-office style storage area network (SAN) using the Dell™ Equallogic™ series of storage arrays. Based on this testing, the Catalyst 3750X meets expectations for use when building SANs that consist no more than four 1 Gb PS Series arrays.

This reference architecture illustrates how to build a medium scale SAN consisting of two Catalyst 3750X switches. It also provides data to show that the Catalyst 3750X can support a small sized SAN using a two-switch solution and can support a maximum of four 1 Gb PS 6x00 arrays with eight 1 Gb hosts in a fully redundant configuration while offering SAN performance for a variety of workloads in smaller storage networking solutions.

2 Objectives

This document provides details on configuring the Catalyst 3750X for use with EqualLogic PS Series storage arrays. The goal of this exercise is not to provide a comprehensive set of possible configurations but to illustrate one possible solution that provides acceptable performance and scalability as validated by testing in our labs.

The test objectives used while testing the Catalyst 3750X configuration are defined below:

- Test the ability of the switch configuration to pass iSCSI traffic as defined by realistic application workloads and server/storage configurations while meeting stringent networking performance parameters.
- Determine the scalability behavior of the switch configuration for a standardized set of I/O workloads and provide sizing guidance in terms of the number of storage arrays and servers that can be supported by a SAN configured with Catalyst 3750X switches.

3 Conclusions

We performed tests that prove the value of the Catalyst 3750X when used to build a SAN with the Dell™ EqualLogic™ series of virtual storage arrays. Based on this testing, we determined that the Catalyst 3750X is an appropriate switch for use when building SANs that consist of one to four 1-Gbps EqualLogic PS Series arrays, such as small business or branch-office solutions.

Link Aggregation Group (LAG) configuration

Our lab tests of this SAN solution yielded the following results:

- With two Catalyst 3750X switches, the SAN scaled easily to support four arrays and eight hosts.
- Sequential Write performance scaled linearly to 100% of the theoretical baseline in terms of throughput as measured at the host.
- Sequential Read Performance scaled linearly to 92% in terms of throughput as measured at the host.
- The Random Read/Write performance in terms of IOPS scaled to 82% of the theoretical baseline on a per-server performance basis.
- TCP retransmissions from arrays, as polled periodically from array counters and SAN Headquarters (SANHQ) were low (< 0.5%) across all test configurations. This is another indicator that there are no bottlenecks or design issues within the switch that limited the ability of the switch to support the reference architecture.

Our test results indicate that the recommended number of arrays is up to four arrays in a two switch configuration. If more arrays are connected the performance may be limited.

Stacked configuration

Our lab tests of this SAN solution yielded the following results:

- With two Catalyst 3750X switches, the SAN scaled easily to support four arrays and eight hosts.
- Sequential Write performance scaled linearly to 100% of the theoretical baseline in terms of throughput as measured at the host.
- Sequential Read Performance scaled linearly to 90% in terms of throughput as measured at the host.
- The Random Read/Write performance in terms of IOPS scaled to 80% of the theoretical baseline on a per-server performance basis.
- TCP retransmissions from arrays, as polled periodically from array counters and SAN Headquarters (SANHQ) were low (< 0.5%) across all test configurations. This is another indicator that there are no bottlenecks or design issues within the switch that limited the ability of the switch to support the reference architecture.

Our test results indicate that the recommended number of arrays is up to four arrays in a two switch configuration. If more arrays are connected, the performance may be limited.

4 Reference architecture

4.1 Reference architecture overview

The reference architecture described in this paper consists of two Catalyst 3750X switches. Connected to these switches are a total of four PS6000XV arrays and eight Dell PowerEdge™ R710 rack servers.

Figure 1 shows the two switch/four array SAN test configuration. The diagram represents both the LAG and Stack options, as described in section 4.4 and section 4.5.

Note: Only the active controller ports are shown.

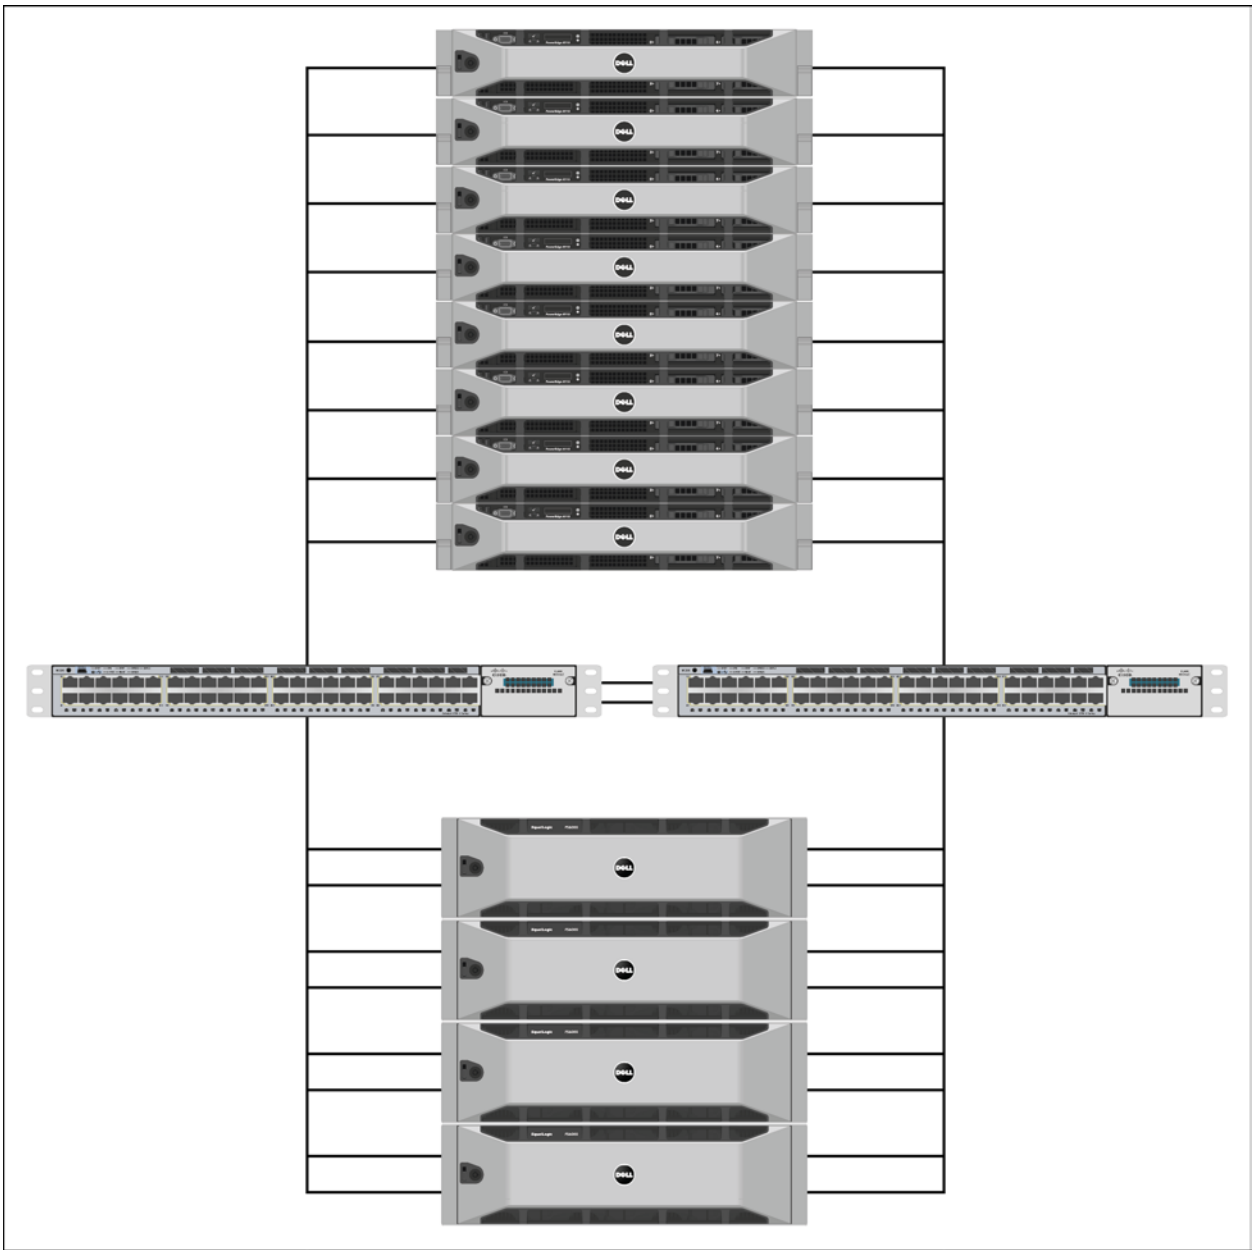

Figure 1 Stack and LAG configurations for the two switch / four array reference configuration

Because our goal is to help you deploy a switch-based SAN easily and quickly, we use a standard host configuration and a standard EqualLogic Group configuration, accept default switch settings wherever possible, and employ all accepted best practice recommendations for EqualLogic SANs.

When developing this reference architecture, we used the following guidelines:

- All hosts have two iSCSI Ethernet ports attached to the SAN.
- All NICs are configured based on default, "out of the box" settings where possible. The exceptions are the use of Jumbo Frames and Flow Control, both of which are enabled for all testing.
- All hosts run Microsoft® Windows Server®.
- We use the EqualLogic Host Integration Toolkit for all hosts. In particular, we use the MPIO Device Specific Module to provide EqualLogic-aware multi-pathing.
- Host connections to the SAN equal the number of active array ports connected to the SAN.
    - Since each PS6x00 series array has four active array ports, for each array in the test configuration two hosts are also connected to the SAN.
- The SAN is on a single, non-routed subnet.
- The SAN uses standard IPv4 addressing.
- Our tests use three pre-defined standardized workloads that reflect various types of real-world SAN utilization.

Note: For more information on EqualLogic SAN design, consult the *EqualLogic Configuration Guide* that can be found at www.delltechcenter.com/page/equallogic+configuration+guide.

4.2 Server configuration

This reference architecture uses eight PowerEdge R710 rack servers.

Table 1 provides the model specifications and configuration settings for each server.

Table 1 Server specifications and configuration settings

PowerEdge R710 Specifications	
BIOS	2.1.15, 9/2/2010
Intel™ 5500-5520 chipset	A05
OS	Microsoft Windows Server® 2008 R2 Enterprise SP1 (Build 7601)
Service Packs and Hotfixes	KB979711, KB976443
Network Interface Card (NIC)	
Model	Intel® Gigabit ET Dual Port Server Adapter
OS Network stack TCP	Autotuninglevel = disabled (default setting 'normal') To disable autotuning run the following command from the CLI: `netsh int tcp set global autotuninglevel=disabled`
iSCSI initiator	Microsoft Windows Server 2008 R2
Intel Gigabit ET Dual Port Server Adapter	Intel® Version: 11.4.7.0 Date: 12/4/2009
MPIO Configuration	
Dell EqualLogic Host Integration Toolkit	Version 3.5.1
Dell EqualLogic MPIO Device Specific Module	Maximum Sessions per Slice: 2 (default) Maximum Sessions per Volume: 6 (default)

4.3 Array configuration

This reference architecture uses four EqualLogic PS 6000XV arrays. All arrays in the Storage Group are the same model and use the same Array Software version (version 5.1.2).

Table 2 Array configuration information

EqualLogic storage	
Array Model	PS 6000XV
Firmware	5.1.2 (R197668)
Enabled performance load balancing in pools	Enable

4.4 Switch configuration option - LAG

The following sections specify the hardware used and settings recommended to configure the switches in this reference architecture when using the LAG configuration.

4.4.1 Switch configuration overview

Table 3 provides an overview of the switch configuration for this SAN.

Table 3 Switch configuration overview (LAG)

Switch settings	
Switch Model	Cisco 3750X
Switch inter-connection (LAG)	Dynamic Link Aggregation Group (LACP – LAG) Flow control enabled on each port channel group
Global Switch Settings	jumbo mtu 9216
Individual Port Settings	flowcontrol on spanning-tree portfast no-storm control unicast Buffer settings (see switch configuration)
Switch Firmware	15.0(1)SE2
Host-Switch Cable Type	CAT6
Array-Switch Cable Type	CAT6
Switch-Switch LAG Cable Type	Cisco SFP Optical Transceiver (SFP-10G-SR); LC-LC Fiber Optic Cable

4.4.2 Global switch settings

Run these commands on both switches.

Configure the Ports

```
switch(config)#int range gig 1/0/1-48
switch(config-if-range)#flowcontrol receive on
switch(config-if-range)#spanning-tree portfast
switch(config-if-range)#no storm-control unicast level
switch(config-if-range)#no shut
switch(config-if-range)#exit
switch(config)#system mtu jumbo 9198
Switch(config)#system mtu routing 9198
```

Configure QOS and optimize buffers for EqualLogic iSCSI use

```
switch(config)#mls qos
```

```
switch(config)#mls qos queue-set output 1 threshold 1 100 100 100 400
switch(config)#mls qos queue-set output 1 threshold 2 3200 100 10 3200
switch(config)#mls qos queue-set output 1 threshold 3 100 100 100 400
switch(config)#mls qos queue-set output 1 threshold 4 100 100 100 400
switch(config)#mls qos queue-set output 1 buffers 4 88 4 4
```

Configuring LAG

```
switch(config)#interface range tenGigabitEthernet 1/1/1-2
switch(config-if)#no shut
switch(config-if)#flowcontrol receive on
switch(config-if)#channel-protocol lacp
switch(config-if)#channel-group 1 mode active
Switch(config-if)#exit
Switch(config)#interface port-channel 1
Switch(config-if)#flowcontrol receive on
Switch(config-if)#switchport trunk allowed vlan all
Switch(config-if)#exit
```

Save the configuration

```
switch(config)#exit
switch#copy run start
switch#reload
```

4.5 Switch configuration option - Stack

The following sections specify the hardware used and settings recommended to configure the switches in this reference architecture when using the Stacking configuration.

4.5.1 Switch configuration overview

Table 4 provides an overview of the switch configuration for this SAN.

Table 4 Switch configuration overview (Stack)

Switch settings	
Switch Model	Cisco 3750X
Global Switch Settings	jumbo mtu 9216
Individual Port Settings	flowcontrol on spanning-tree portfast no-storm control unicast Buffer settings (see switch configuration)
Switch Firmware	15.0(1)SE2
Host-Switch Cable Type	CAT6
Array-Switch Cable Type	CAT6
Switch-Switch Stack Cable Type	Cisco Stacking cable

4.5.2 Global switch settings

Run these commands on both switches.

Set up the ports

```
switch(config)#int range gig 1/0/1-48, gig 2/0/1-48
switch(config-if-range)#flowcontrol receive on
switch(config-if-range)#spanning-tree portfast
switch(config-if-range)#no storm-control unicast level
switch(config-if-range)#no shut
switch(config-if-range)#exit
switch(config)#system mtu jumbo 9198
switch(config)#system mtu routing 9198
```

Configuring QOS and optimize buffers for EQL iSCSI use

```
switch(config)#mls qos
switch(config)#mls qos queue-set output 1 threshold 1 100 100 100 400
switch(config)#mls qos queue-set output 1 threshold 2 3200 100 10 3200
switch(config)#mls qos queue-set output 1 threshold 3 100 100 100 400
```

```
switch(config)#mls qos queue-set output 1 threshold 4 100 100 100 400
switch(config)#mls qos queue-set output 1 buffers 4 88 4 4
```

Save the configuration

```
switch(config)#exit
switch#copy run start
switch#reload
```

Appendix A Switch port mappings

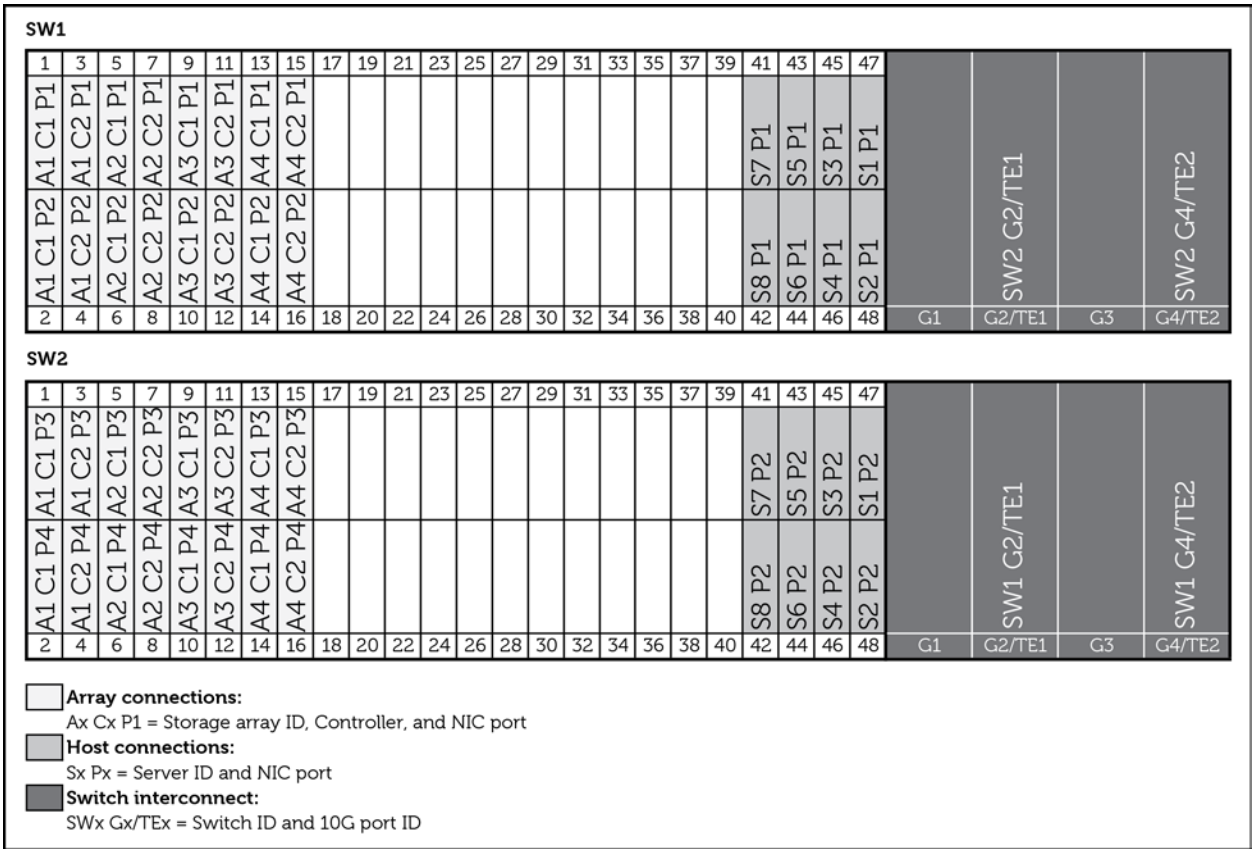

Figure 2 Switch configuration option – LAG: Switch port mappings for the two switch configuration

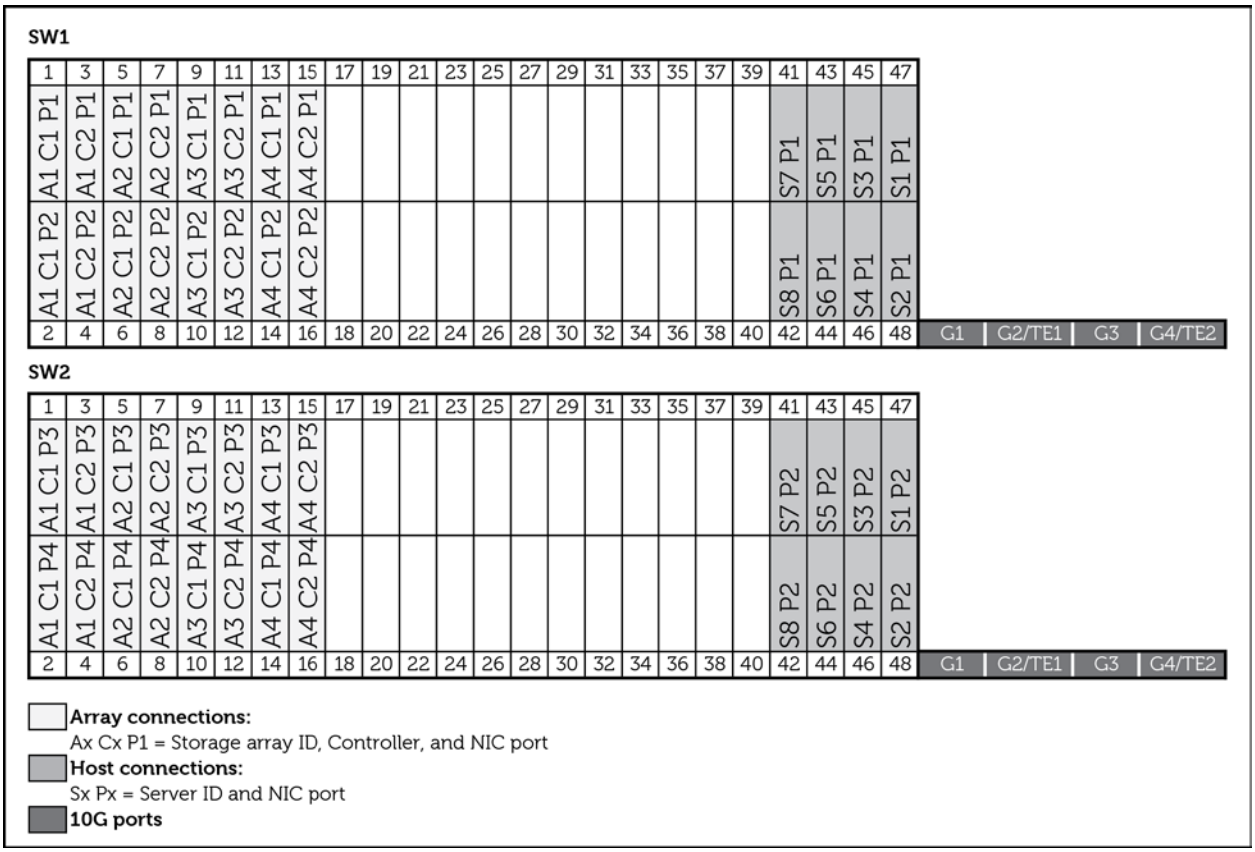

Figure 3 Switch configuration option – Stack: Switch port mappings for the two switch configuration

THIS WHITE PAPER 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.