



Best Practice of Server Virtualization Using Qsan SAN Storage System

**F300Q / F400Q / F600Q Series
P300Q / P400Q / P500Q / P600Q Series**

**Version 1.0
July 2011**

Copyright

Copyright@2011, Qsan Technology, Inc. All rights reserved. No part of this document may be reproduced or transmitted without written permission from Qsan Technology, Inc.

Trademarks

All products and trade names used in this manual are trademarks or registered trademarks of their respective companies.

Qsan Technology, Inc.

2F., No.23, Lane 583, Ruiguang Rd.
Neihu Dist., Taipei 114
Taiwan, R.O.C.

Tel: +886-2-7720-2118
Fax: +886-2-7720-0295

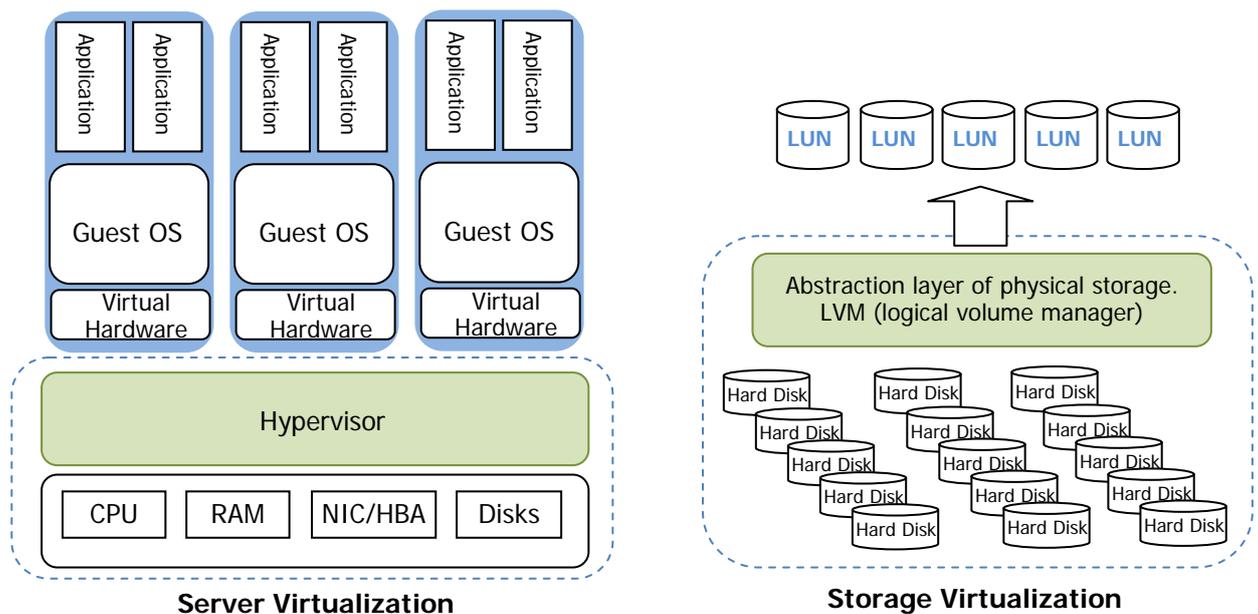
Email: sales@Qsan.com.tw
Website: www.QsanTechnology.com

Abstract

You can't take full advantage of the power of server virtualization without the involvement of SAN (Storage Area Network) storages. It's the most important ingredient in making a successful, efficient, and highly available clustering recipe using server virtualization software. In this article, we will outline the keystones and summarize how to achieve high availability design using SAN storage from a structural point of view instead of step by step manual instructions. A great IT infrastructure starts with a careful and thoughtful evaluation and planning work. Qsan SAN products will assure you a cost-effective and successful deployment of server virtualization. Here we assume our readers are familiar with the basic operations of VMware ESX server and VirtualCenter.

Preface

"Virtualization" is a buzz word these days. There are many different kinds of virtualization in computer technology. Simply put, virtualization means abstraction. The virtualization solution abstracts the underlying details of complexity of whatever it is virtualizing. The goal of virtualization is to make things simpler to manage. There are two major kinds of virtualization - storage virtualization and server virtualization. It depends on where the virtualization happens along the data flow from host computers to storage devices. Here we only focus on the most popular one - server virtualization. Our topic here is to show you that IP SAN storage plays such an important role in realizing the most value-added functions of server virtualization and how you can make the best of it.



Server Virtualization - Use Less to Do More

Server virtualization has been around for a long time. It can be categorized into hardware-based server virtualization and software-based one. Hardware-based server virtualization refers to powerful mainframe computers like HP blade servers, IBM and SUN mainframes. It is known as logical partitions (LPARs). The mainframe computer partitions its physical resources into logical partitions to run different jobs at the same time, enabling better utilization of its very expensive hardware. These domain-class monsters use high-speed backplane, share common I/O ports and memory to achieve massive scalability in less footprint to replace less capable standalone traditional servers.

Software-based server virtualization on the other hand abstracts a single physical device and turns it into multiple virtual devices. It creates an abstraction layer between the physical server hardware and the virtual instances (virtual machines) that run on top of that hardware. There are 3 major players in this area, which are VMware, Citrix, and Microsoft Hyper-V.

Benefits of software-based server virtualization:

- Server consolidation : reduce the number of physical servers, reduce hardware costs
- Reduce floor space, less heat and cooling requirements, less power consumption
- Lower cabling and infrastructure costs by sharing virtual connections
- Better and efficient utilization of hardware components
- Simplified server management. Swap, add, or change hardware components in just a few mouse clicks.
- High availability and disaster recovery become cost-effective and simple.

Network Storage is the Keystone and Foundation

Instead of increasing hardware utilization rate, the most valuable part of server virtualization is its ability to provide high availability and fault tolerance functions. In order to achieve this goal, server virtualization needs its best friend - network storage. The idea relies on that all the hosts can access the same storage volume where all the virtual machine files are stored. So when any host is offline, the backup hosts can take over the virtual machines running on the failed host instantly and non-disruptively. Without network storage, none of these will be possible. This truly sets SAN storage apart from the

traditional DAS (Direct Attached Storage) storage. Because when you use the traditional direct-attached RAID storage, if the server goes down, the whole system goes down. It doesn't matter if you have RAID protection. What matters is networking capability of your storage devices, which makes Qsan iSCSI products (P300Q, P400Q, P500Q, P600Q) and Fibre Channel products (F300Q, F400Q, F600Q) a perfect candidate for this type of application.

The idea behind creating highly available server clustering is simple. By creating a virtual server farm across multiple powerful physical servers, and storing all of your virtual machines on a centralized SAN target, which is fully redundant and does not have single point of failure. You can ensure that the failure of any given piece of physical hardware will not bring down your virtual environment.

In choosing your network storage, SAN and NAS are two common options. SAN generally has better performance and more advanced functions over NAS. It's the recommended solution in server virtualization. There are two popular protocols in SAN storage - Fibre Channel and iSCSI. Generally speaking, Fibre Channel solution is much more expensive than iSCSI solution. Considering ubiquity, cost efficiency, and technology familiarity, iSCSI would be the choice for most small and medium business companies. About Qsan Fibre Channel SAN storage products, you can choose from 4Gb or 8Gb FC host interface. Fully redundant and dual controller design comes with complete RAID levels of data protection. Your data are always safe with Qsan. The other more affordable and popular choice is iSCSI. Qsan is an iSCSI expert. A complete product lineup includes P300Q, P400Q, P500Q, and P600Q, which offers host interfaces of 1GbE and 10GbE with the number of iSCSI ports ranging from 8 to 12 in dual controller mode.

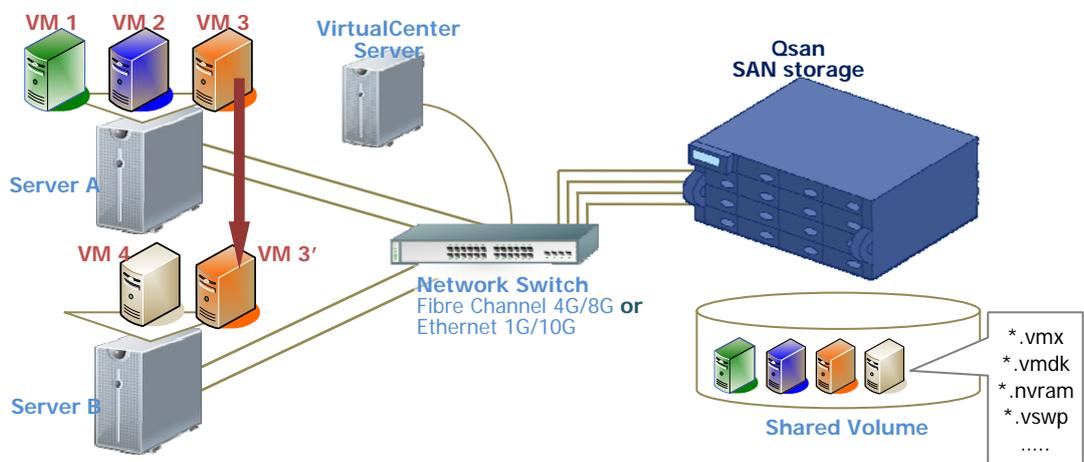
Model	Host Interface	Supported HDDs	VMware HA, FT	Citrix HA, FT	Fully Redundant design
P300Q	8x 1Gb iSCSI	3Gb, SAS/SATA/SSD	Yes	Yes	Yes
P400Q	12x 1Gb iSCSI	6Gb, SAS/SATA/SSD	Yes	Yes	Yes
P500Q	4x 10Gb iSCSI	3Gb, SAS/SATA/SSD	Yes	Yes	Yes

Model	Host Interface	Supported HDDs	VMware HA, FT	Citrix HA, FT	Fully Redundant design
F300Q	8x 4Gb FC	3Gb, SATA/SAS/SSD	Yes	Yes	Yes
F400Q	8x 4Gb FC 2x 1Gb iSCSI	3Gb, SATA/SAS/SSD	Yes	Yes	Yes
F600Q	8x 8Gb FC	6Gb, SAS/SATA/SSD	Yes	Yes	Yes

All three server virtualization products offer high availability and fault tolerance clustering functions. However, when it comes to product maturity, popularity, and completeness, people talk about VMware more often. So we are using VMware solution as our reference in this article.

	VMware	Citrix	Hyper-V
High Availability	HA, DRS, vMotion	XenMotion	Live Migration
Supported Version	ESX 3.5 or above VMware Infrastructure Enterprise 3 or above	Enterprise or above	R2

The basic conceptual structure will look like this.



Server A and server B make up a cluster using VMware HA. A virtual machine is nothing but a set of files such as *.vmx, *.vmdk, *.nvr and *.vswp instead of physical hardware components. Store all the virtual machines' (VM1 ~ VM4) files at Qsan storage volume that is shared by both servers, which means both server A and server B have access to all virtual machines' files.

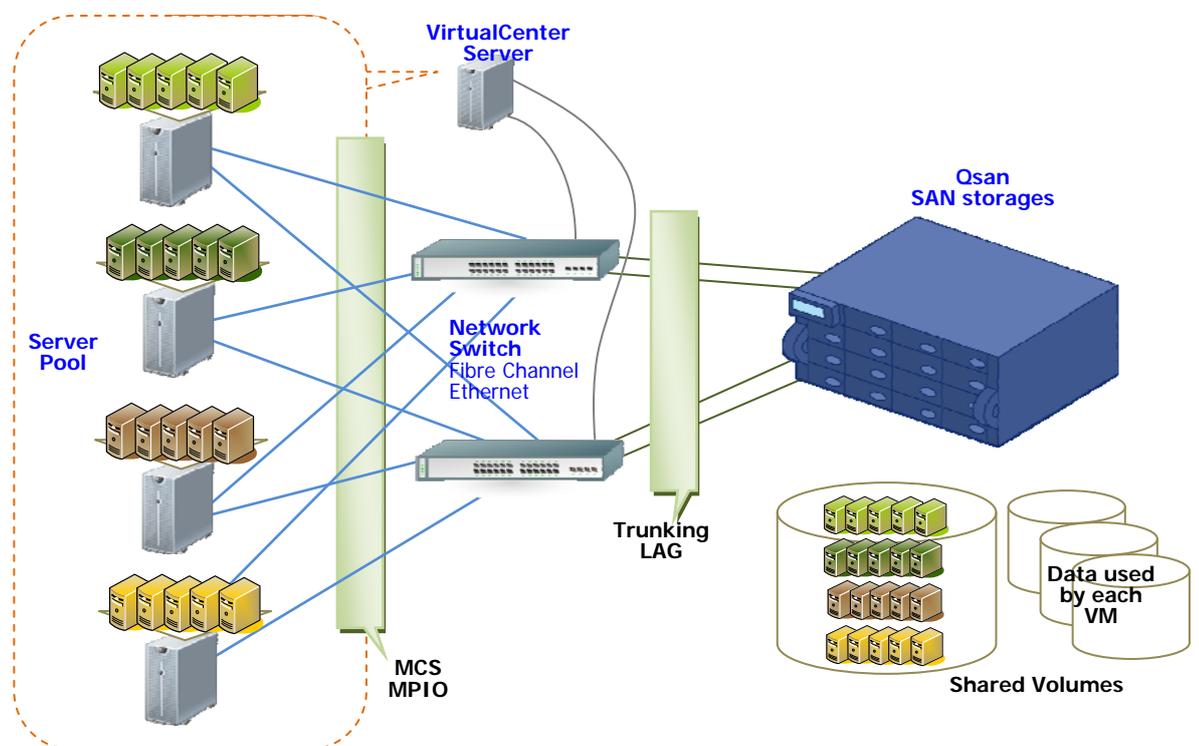
Suppose that we only set the rules for VM3. When Server A's hardware is failing, VirtualCenter detects this and have Server B launch VM3 on itself. The running status of VM3 is migrating from server A to server B. All the transition happens in just a few seconds if you have a fast storage system to support this.

You can make the network switch redundant by adding another one. The VirtualCenter server itself can be implemented as VM1 or VM2 to be protected by high availability. There you go. With no extra cabling and no extra hardware equipments, if you use Qsan

SAN products as your storage solution, you can get the extra benefits of implementing a very cost-effective fault tolerant system.

How to Arrange Qsan SAN Storages to Make It Work?

On the host side, you use the cluster function provided by VMware to achieve high availability and fault tolerance of virtual machines. On the storage side, high availability and fault tolerance are protected by Qsan dual active controller design. Between the host and storage, MPIO and link aggregation functions are provided by Qsan storages to ensure redundant data paths and aggregate performance are in place, which provides both load balancing and failover. So your end-to-end data flows are well protected. Below we show you how to design the whole systems using Qsan iSCSI and Fibre Channel SANs.



In the graph above, we show the actual building blocks that constitutes a fault tolerant design in production scale. If the server pool has more than 2 servers, DRS can be enabled for better hardware utilization ratio. To add more security to the network switch part, which most people will do, redundant switch design is preferred. When one switch goes down, you have the second switch for backup solution and all connections and data routes can go through the second switch and remain intact. Extra bonus is that the redundant path is not just skin deep, improved performance can be achieved by

aggregating multiple connections. Again, both Qsan iSCSI SAN and Fibre Channel SAN products can fit into this structure perfectly.

On the storage side, Qsan SAN storages provide shared volumes to keep all virtual machine files safe and secure. The shared volume can be accessible to all servers in the server pool. Advanced features such as trunking and link aggregation are provided by Qsan SAN storages to achieve redundant connections and better performance. As we mentioned earlier, SAN storage makes all these possible by shared volume design and networking capability. High availability and vMotion function are provided by VirtualCenter server, which manages and monitors the server pool. VirtualCenter can run on a standalone physical server or on a virtual machine.

If you use fiber cables for the network connections, servers can even be placed at different buildings or sites within a few kilometers. This gives you more flexibility in designing your IT infrastructure.

A Checklist to Setup a High Available and Fault Tolerant Application

- Qsan iSCSI SAN or Fibre Channel SAN.
Performance requirement : throughput at least 200MB/sec
- Each server has 2 identical HBA cards. Each HBA has at least 4~8x 1gigabit Ethernet ports or at least 2x 4gigabit Fibre Channel ports.
- VMware infrastructure Enterprise edition version 3 or above (with VirtualCenter).
- A dedicated server or virtual machine running VirtualCenter Server and VirtualCenter Database
- Gigabit or 10Gigabit Ethernet switches for iSCSI SAN storage. or
4G or 8G Fibre Channel switches for FC SAN storage.
- Connections to the same physical network.
- Hosts' CPUs are compatible to run each virtual machine.
- Fine tune the options and settings
 - Arrange capacity on each host such that when losing one ESX host/server, the server pool still has enough capacity to run every virtual machine normally.
 - HA settings, placement constraints, virtual switch port group, migration priority

Once you have all the elements in place, a successful deployment is just a few mouse clicks away. You can enjoy the benefits and reliability brought to you by server virtualization.

Conclusion

Everyone managing a storage system cares about how safe it is and how to make it safer. In the past, high availability and fault tolerance applications can be achieved by very expensive hardware combinations that only large enterprises can afford. Server virtualization promises to make high availability and disaster recovery more straightforward and cost-effective than ever. It can recover from server failures by automatically restarting virtual machines on alternate servers, which makes practically any application highly available and fault-tolerant. But in order to make this work, you will need Qsan SAN storage subsystems to provide high speed data transmission and upmost data protection.

Let Qsan SAN storage solutions help you increase storage utilization and ROI as well as drive down unnecessary costs. It's your best choice to optimize your IT budget and make every penny work. Should you have any technical queries, please contact support@qsan.com.tw. If you are interested in receiving more product information, please contact sales@qsan.com.tw.

Applied To

- F300Q / P300Q / P500Q Series: FW 2.1.0
- F400Q / F600Q / P400Q / P600Q Series: FW 1.0.0

References

- Qsan White Paper
[QWP200917-P300H-Connect_P300H_with_iSCSI_initiator_in_ESX4.0.pdf](#)
- Qsan White Paper
[QWP201002-P210C-Introduction_to_MPIO_MCS_Trunking_and_LACP.pdf](#)
- VMware Training
<http://mylearn.vmware.com/mgrreg/>
- VMware Server
<http://www.vmware.com/products/server/overview.html>