

## USING VIRTUAL IRON 4.5 WITH SAN STORAGE

---

### About Virtual Iron

---

Virtual Iron is an enterprise-class server virtualization technology that fully leverages the virtualization capabilities of today's modern x86 processor chips as well as industry standards. Virtual Iron's technology can be used to address many of the key server management challenges within a data center including:

- Efficient utilization of server resources
- Rapid provisioning and deployment
- High availability and recovery
- Workload management and
- Cost containment

Virtual Iron allows a physical server to be logically partitioned to run multiple instances of virtual servers that each run their own operating system image and are allocated a specific number of CPUs (which are shared amongst the other virtual servers) and a fixed amount of dedicated memory. In addition, each virtual server has access to one or more virtual Ethernet network interfaces and storage (disks). The entire virtual environment is managed by Virtual Iron's VI-Center. VI-Center is used to perform all configuration and management functions in the virtual environment.

For more detailed information on Virtual Iron's architecture please visit the web site at <http://www.virtualiron.com>.

### Virtual Iron Storage Configuration

---

#### Physical Disk Connectivity

Physical disk storage must first be made accessible to the managed node before it can then be provisioned to the virtual servers. The storage provisioning process for the managed node is generally the same as the provisioning process for any other physical server— all the same provisioning steps are required (storage LUN configuration, zoning, LUN masking, discovery, etc.).

There are three storage options:

1. *Fibre Channel (FC) SAN*. In this case, one or more FC host bus adapters (HBAs) must be installed in the managed node. The managed node can then connect to any supported storage subsystem accessible within the FC fabric. Note that all the normal Fibre Channel storage provisioning tasks must be performed to make the desired LUNs available to the managed node.
2. *iSCSI*. In this case the managed node will access Ethernet-attached storage via a NIC. Note that all the normal iSCSI storage provision tasks must be performed to make the desired LUNs available to the managed node.

3. *Local SATA, SAS or parallel SCSI disks.* In this case the disk is accessed via the appropriate SCSI adapter installed in the managed node. Local storage is not covered in this paper.

LUNs presented to the managed node may be physical disks (e.g. in the case of locally attached SATA, SAS or SCSI disks) or abstracted LUNs from a storage subsystem controller (e.g. a “slice” of group of disks in a RAID-5 configuration). In either case, the VI-Center and the managed node automatically discover the LUNs. Disks are then made accessible to virtual servers in one of two ways: As a logical disk or as a raw disk.

## Logical Disks

A logical disk represents a specific quantity of capacity (e.g. 10GB) within a disk group. When a logical disk is created via VI-Center, behind the scenes a Linux logical volume is created and presented to the virtual server as a SCSI disk drive. Note that there is no file system layered between the logical volumes on the service partition and the disk LUNs on the virtual servers.

Compared to raw disks, logical disks have the following advantages and disadvantages:

- Advantage: Easier to manage and allocate storage resources without SAN changes. For example, in the physical world, if 10 servers each required a 10GB LUN, the array storage would have to be carved up into 10 separate LUNs and then the storage provisioning process performed for each server (zoning and LUN masking). However, with 10 virtual servers, the provisioning of the storage only has to take place through VI-Center (one time instead of 10 times). Then logical disks can be created and assigned to each virtual server.
- Disadvantage: Storage subsystem tools can't be used on individual virtual disks since the SCSI pass-through commands are not supported.

## Raw Disks

As a raw disk, a LUN on the service partition is presented completely and directly to the virtual server. In this case, the service partition's logical volume manager is not used to access the disk.

Compared to logical disks, raw disks have the following advantages and disadvantages:

- Advantage: Storage subsystem tools (e.g. snapshot, replicate, etc.) can be used on individual virtual server disks (LUNs).
- Disadvantage: VI-Center can't be used to perform useful functions like snapshot, clone, move, import and export.

## SAN Storage Overview

---

It is important to plan your storage configuration in advance of deploying virtual infrastructure. Here are some guidelines to keep in mind:

1. Take care when adding, removing, and resizing LUNs as it may require a physical server reboot. Do not resize LUNs that are used as part of Logical disks; instead, create a new LUN and add it to the disk group.

2. Test your configuration, especially failover, in a test environment before rolling into production. If your array firmware is at a different release number than we have tested, confirm whether there are any differences. You may need to make changes to the multipath configuration files found in VI-Center.
3. Plan the size and type of storage that you are using by workload. For example:
  - Boot volumes can typically be on larger drives as most operating systems have minimal I/O activity on the boot disk, but some of that I/O is memory paging, which is sensitive to response times.
  - Applications can be on larger, slower drives (e.g. RAID 5) unless they perform a lot of I/O. Write-intensive workloads should use RAID 10 on medium to fast drives. Ensure that log files are on different physical drives than the data they are protecting.
  - Infrastructure servers such as DNS tend to have low I/O needs. These servers can have larger, slower drives.
4. If using storage server features such as cloning and snapshots, use Raw disks.
5. While it may be tempting to create a very large LUN when using logical disks, this can be detrimental to performance as each VM queues I/Os to the same disks. We recommend logical disk LUNs do not exceed 300-500GB.

## Virtual Iron SAN Testing

---

For each array that has been certified, there is a chart that describes the tests that have been performed:

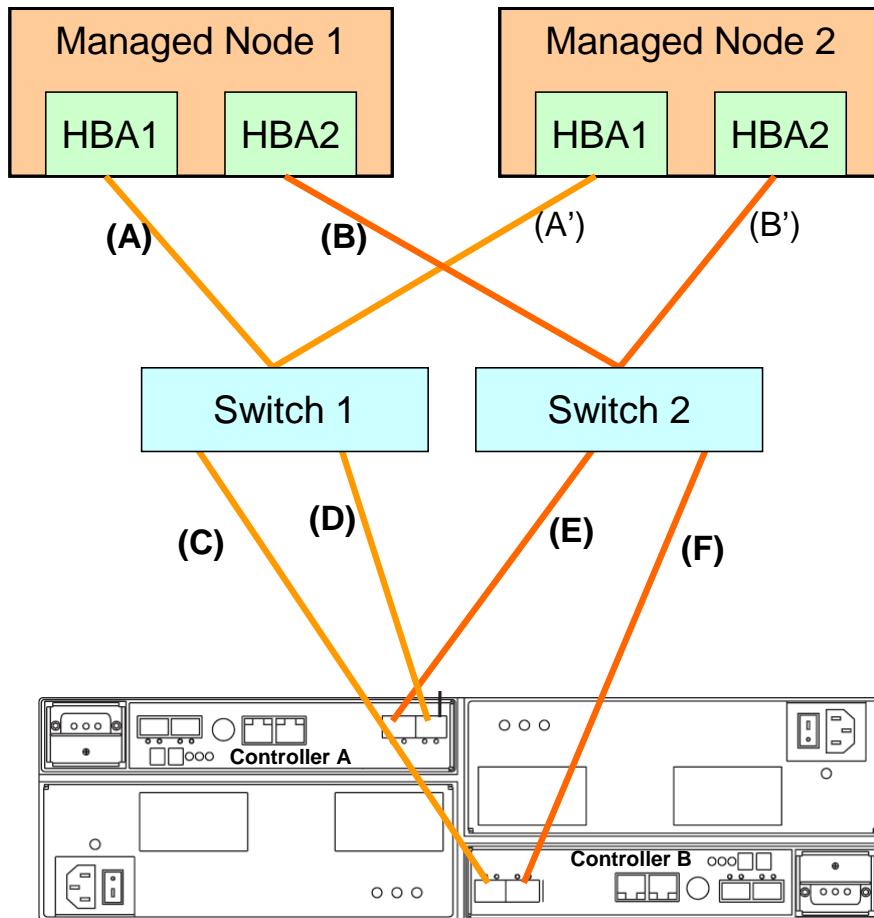
**Basic Connectivity** — The ability for Virtual Iron to recognize and interoperate with the storage array. This configuration alone does *not* support failover. It assumes a single path to a single controller.

**Node Port Failover** — The ability for the Virtual Iron node's NIC or HBA to be configured with more than one NIC or HBA ports to one or more switches. Virtual server I/O is robust to port failure and switch failure.

**Storage Port Failover** — The storage server has multiple storage ports and the configuration is robust to storage port failure.

Note: Some storage arrays manage their ports transparently to Virtual Iron. When using one of these systems, Virtual Iron communicates to the virtual port address. In the event that communication with a port is lost, Virtual Iron automatically reconnects with the virtual port and the storage server redirects to an active, usable port. This reconnection process can delay virtual machine I/O for up to sixty seconds while the failover takes place. Virtual Iron recommends setting the Windows registry key `KEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk\TimeoutValue` to 120 seconds to avoid virtual machine disk errors.

A diagram of a typical configuration is shown below:



## Compellent

---

	iSCSI	Fibre Channel
Basic Connectivity	✓	✓
Node Port Failover		
Storage Port Failover		

## DataCore SANmelody

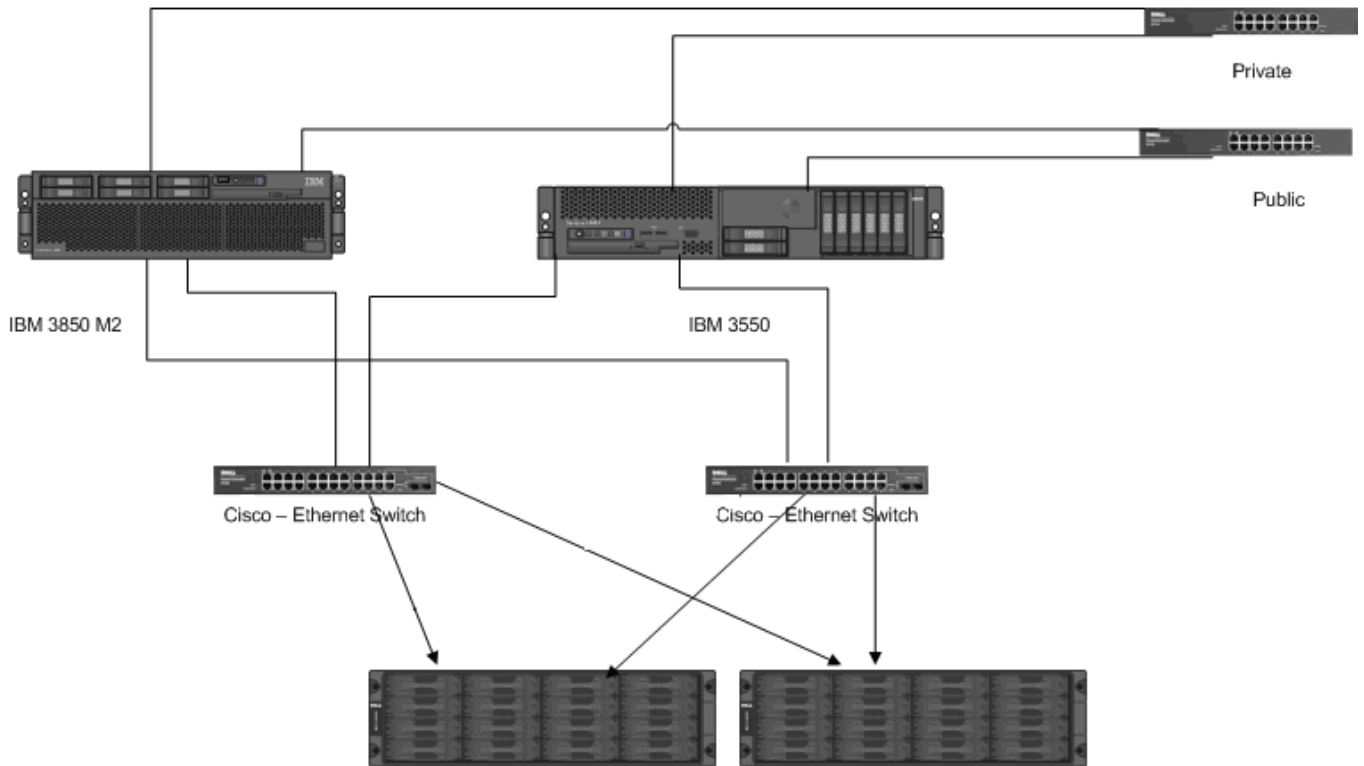
---

	iSCSI	Fibre Channel
Basic Connectivity	✓	✓
Node Port Failover		
Storage Port Failover		

## Dell MD3000i

---

The following is a diagram of the test configuration:



	iSCSI
Basic Connectivity	✓
Node Port Failover	
Storage Port Failover	

## EMC CX150i / Dell AX150i

---

	iSCSI
Basic Connectivity	✓
Node Port Failover	
Storage Port Failover	

## Equallogic

---

	iSCSI
Basic Connectivity	✓
Node Port Failover	
Storage Port Failover	

## FalconStor IPStor

---


	iSCSI
Basic Connectivity	✓
Node Port Failover	
Storage Port Failover	

## HP MSA 1510i

---

Model name	Controller Firmware Version
MSA 1510i	2.00b803 *  Note: Older firmware versions do not work reliably.

The following is a diagram of the test configuration:




	iSCSI
Basic Connectivity	✓
Node Port Failover	
Storage Port Failover	

## HP MSA 2012i

---

The following is a diagram of the test configuration:



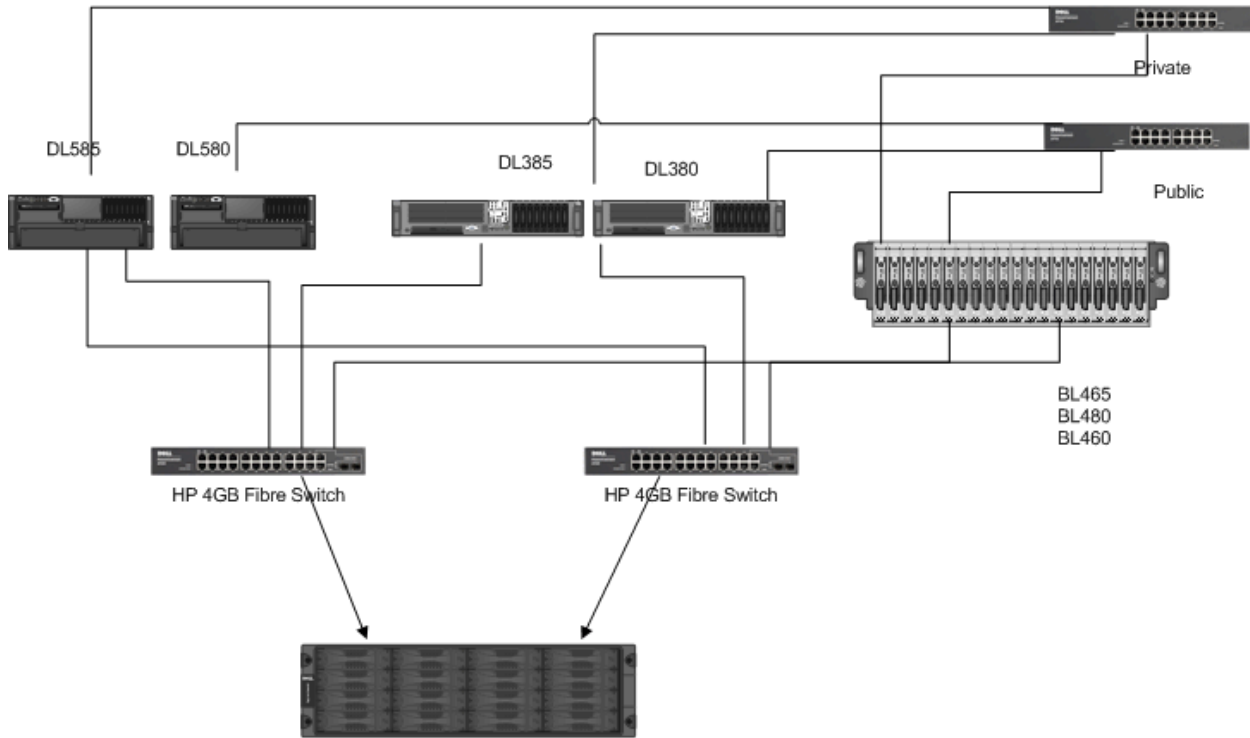


	iSCSI
Basic Connectivity	✓
Node Port Failover	
Storage Port Failover	

## HP EVA 4400

Model name	Model Number	Model Version	Controller Firmware Version
EVA 4400	HSV300		09000000

The following is a diagram of the test configuration:




	Fibre Channel
Basic Connectivity	✓
Node Port Failover	✓
Storage Port Failover	✓

## IBM DS3300

Model name	Model Number	Model Version	Microcode Level
DS3300	1726		06.70.22.00

The following is a diagram of the test configuration:




	iSCSI
Basic Connectivity	✓
Node Port Failover	
Storage Port Failover	

### IBM DS4100, DS4300, DS4700, DS4800, DS6000 Series, DS8000 Series (8100, 8300)

Model name	Model Number	Model Version	Microcode Level
DS4100	1724	1RU	06.12.23.00
DS4300	1742	1RU	06.12.27.00
DS4700	1814		06.16.88.00
DS4800	1850		06.16.88.00
DS6000	1750	511	5.2.2.393
DS8100	2107	921	5.2.420.531

The following is a diagram of the test configuration:




	FC
Basic Connectivity	✓
Node Port Failover	✓
Storage Port Failover	✓

## Network Appliance FAS Series and V-Series

---

Tested with Data ONTAP Release 7.2.5.1



	iSCSI	Fibre Channel
Basic Connectivity	✓	✓
Node Port Failover		
Storage Port Failover		

## StoreVault

---

	iSCSI
Basic Connectivity	✓
Node Port Failover	
Storage Port Failover	

## Microsoft Windows Storage Server 2003

---

	iSCSI
Basic Connectivity	✓
Node Port Failover	
Storage Port Failover	

## Fibre Channel Host Bust Adapters

---

Virtual Iron supports the following host bus adapters:

<b>Emulex</b>	<b>LSI</b>	<b>Qlogic</b>
LPe11000	LSI7204EP	QLA200
LPe11002		QLA2200F
LPe1105		QLA2340
LPe1150		QLA2344
LP11000		QLA2462
LP11002		QLE2360
VLP1150		QLE2460
LP10000ExDC		QLE2462
LP1050Ex		
LP10000		
LP10000DC		
LP1050		
LP1050DC		
LP9802		
LP9802DC		
LP982		

LP9000		
--------	--	--