



Technical Report

PowerVM Virtualization and NetApp Storage: Best Practices

Bobby Oommen, NetApp Inc.

December 2009 | TR-3806

EXECUTIVE SUMMARY

IBM® Power Systems™ servers coupled with PowerVM™ technology enable clients to build a dynamic infrastructure, reducing costs, managing risk, and improving service levels. By using NetApp® storage solutions as the underlying storage, customers can make the best use of NetApp virtualization technologies, including virtual data center backup and restore with simple and flexible business continuance, and minimal-cost cloning of LPARs. This technical report describes best practices for setting up IBM Power Systems using NetApp products for the underlying storage system.

TABLE OF CONTENTS

1	INTRODUCTION	3
2	BUSINESS CASE	3
3	BUSINESS SOLUTION	3
4	IBM AIX VIRTUALIZATION COMPONENTS	4
4.1	POWERVM.....	4
4.2	POWER HYPERVISOR.....	4
4.3	LOGICAL PARTITIONS (LPAR)	4
4.4	VIRTUAL I/O SERVER (VIOS)	4
4.5	VIRTUAL ETHERNET ADAPTER (VEA).....	4
4.6	SHARED ETHERNET ADAPTER (SEA)	5
4.7	HOST ETHERNET ADAPTERS (HEA).....	5
4.8	N_PORT ID VIRTUALIZATION (NPIV).....	5
6	NETAPP STORAGE AND VIRTUALIZATION	6
5.1	AGGREGATE AND FLEXIBLE VOLUME (FLEXVOL).....	6
5.2	THIN PROVISIONING	6
5.3	FLEXCLONE.....	6
6	HARDWARE CONFIGURATION	7
7	ENVIRONMENT CONFIGURATION	8
7.1	VIRTUAL I/O SERVER.....	8
7.2	NETWORK CONNECTIVITY	9
7.3	AIX LPAR (VIO CLIENT).....	9
7.4	AIX TUNING PARAMETERS	10
7.5	FIBRE CHANNEL ADAPTERS	10
7.6	SAN SWITCH	11
7.7	NETAPP STORAGE SYSTEM	11
8	WORKLOAD TEST	13
9	CONCLUSION	14
10	ACKNOWLEDGEMENT	14
11	APPENDIX	15

1 INTRODUCTION

PowerVM virtualization is a combination of software and hardware components that creates virtual machines by dividing the physical resources such as CPUs, memory, storage, and optical devices. Virtualization helps you regain control of your IT infrastructure by allowing maximum resource utilization and making the computing resources more manageable. With virtualization, you have a logical rather than a physical view of computing power, storage capacity, and other resources. This technical report briefly introduces virtualization using IBM PowerVM and also discusses the essential best practices for using NetApp as the underlying storage system.

2 BUSINESS CASE

Shrinking IT budgets, shaky economic conditions, and explosive data growth are forcing businesses to do more with less and to adopt consolidation and virtualization strategies in their data centers. Virtualization technologies save hardware and software costs, and they can also reduce energy and data center footprint costs. The benefits of virtualization can vary from enterprise to enterprise, depending on the business infrastructure. The business cases for virtualization are:

- **Cost savings.** Virtualization allows consolidation of resources that can lower hardware acquisition costs by purchasing fewer systems. This helps to shrink the infrastructure footprint by decreasing rack space and cutting back on cooling and power.
- **Cost avoidance.** Costs can be avoided by enabling automatic failover to improve the disaster recovery process and business continuity operations. It also helps businesses to grow faster with less staff and increased efficiency.
- **Maximum resource utilization.** Server and storage resources can be allocated and released to keep up with changing business needs. Resource management overhead can be reduced by consolidating multiple servers. The PowerVM virtualization layer separates the hardware resources from the operating system environment, so a lost partition can be recovered on a machine without any hardware configuration, because all of the configuration is available in the management console.

3 BUSINESS SOLUTION

Virtualization is about improving resource utilization and response times of the IT department to meet increasing business demands. Here are the three main business benefits:

- **Improve total cost of ownership.** Virtualization can offer a rapid return on investment through decreased management costs and increased resource utilization. You can enhance system availability and help lower cost and complexity by virtualizing resources such as CPU, memory, Ethernet, and FCP devices at the server level, and by using NetApp technologies such as thin provisioning and FlexClone® at the storage level. With virtualization, it is possible to capture short-term market opportunities while minimizing capital investment.
- **Increase flexibility and simplicity.** By using virtualization to consolidate resources, customers you can increase IT resource utilization and eliminate infrastructure sprawl. Infrastructure provisioning can be done centrally at the enterprise level, enabling dynamic resource allocation based on business demands. Virtualization allows you to add new services dynamically without going through infrastructure complexities. New services can be added quickly, within minutes rather than weeks, enabling you to respond to changes in business requirements. Virtual infrastructure services use virtual resources that can be deleted dynamically and easily without affecting other business functions and services, allowing the freed-up resources to go back to the shared pool of resources for redeployment.
- **Enable access through shared infrastructure.** Virtualization can provide a resilient foundation and shared infrastructure that enable better access to resources. Resources under virtualization are considered a pool, and they can be shared by all applications. Virtualization can help an IT department to secure and isolate application workloads for easier replication and recovery. The IT department can exploit this added resiliency to maintain a highly available infrastructure while performing planned maintenance, and to configure low-cost disaster-recovery solutions.

4 IBM AIX VIRTUALIZATION COMPONENTS

4.1 POWERVM

PowerVM, formerly known as Advanced Power Virtualization (APV), is a combination of software and hardware that delivers industry-leading virtualization on IBM POWER5™ and POWER6® servers. It is used today by thousands of customers for large-scale mission-critical production environments. The introduction of PowerVM made virtualization possible not just at the CPU and memory level; it also became possible to share Ethernet adapters and to virtually slice disk resources for provisioning to the end user. As a result, IBM System p® can now facilitate server consolidation, reduce costs, provide a whole new level of virtualized redundancy, and adjust and optimize all forms of system resources and application workload capacities to quickly respond to changing business demands. PowerVM technologies include Micro-Partitioning™, Logical Partitioning (LPAR), Virtual I/O (VIO), POWER Hypervisor™, Live Partition Mobility (LPM), and so on.

4.2 POWER HYPERVISOR

The POWER Hypervisor is the foundation for virtualization on a System p server. The hypervisor orchestrates and manages system virtualization, including creating logical partitions and dynamically moving resources across multiple operating environments. The hypervisor provides the ability to divide physical system resources into isolated LPARs. Each LPAR operates like an independent system running its own operating environment. The hypervisor can assign dedicated processors, I/O, and memory, which you can dynamically reconfigure as needed to each LPAR. The hypervisor can also assign shared processors to each logical partition by using its Micro-Partitioning feature. Unknown to the logical partitions, the hypervisor creates a shared processor pool from which it allocates virtual processors to the logical partitions as needed. In other words, the hypervisor creates virtual processors so that logical partitions can share the physical processors while running independent operating environments.

4.3 LOGICAL PARTITIONS (LPAR)

Logical partitioning is the ability to logically slice up a single system's CPU, memory, and other resources to create multiple and separate servers. IBM introduced LPAR with the release of AIX® 5.1 and POWER4™ technology as a way to provide greater flexibility and better use of large systems. Using LPARs, systems can run both AIX and Linux® in separate partitions. Each LPAR has its own allocation of CPU, memory, and I/O devices. This type of partitioning is done at the hypervisor level, not at the physical resource level. Thus with IBM System p5® servers, it is possible to create LPARs using Micro-Partitioning—assigning less than one physical CPU per LPAR.

4.4 VIRTUAL I/O SERVER (VIOS)

VIOS is part of the IBM PowerVM Virtualization hardware feature. This is a special-purpose partition that provides virtual I/O resources to client partitions. The VIOS owns hardware adapters like SCSI disks, Fibre Channel disks, and Ethernet or CD/DVD optical devices, but allows other LPARs to access them or to be part of them. It is possible to share a physical adapter by assigning it to multiple client partitions, which minimizes the number of physical adapters required for individual clients. This special-purpose LPAR with its resources is called the VIO Server, and the other LPARs using it are called VIO clients. Thus the VIOS can reduce costs by eliminating the need for dedicated network and disk adapters.

4.5 VIRTUAL ETHERNET ADAPTER (VEA)

The VEA provides an internal communication path between partitions on the same server. Each VEA appears as a regular Ethernet adapter to the owning LPAR and can be configured with its own IP address. VEAs communicate with each other through the hypervisor. VLANs can be used on the VEAs to segregate Ethernet traffic to specific partitions. VLAN IDs are configured on the Hardware Management Console (HMC) as part of each partition's profile. A good application of the VEA is application-partition-to-database-partition communications within the same server. Using the VEA, customers can reduce latency and improve throughput. In addition, implementing VEAs on the server can help reduce the physical Ethernet

adapter requirements for each partition and significantly reduce the total quantity of Ethernet adapters and associated I/O drawer space at the server level.

4.6 SHARED ETHERNET ADAPTER (SEA)

The shared Ethernet adapter hosted in the VIOS acts as a layer-2 bridge between the internal virtual and external physical networks. The SEA enables partitions to communicate outside the system without having to dedicate a physical I/O slot and a physical network adapter to a client partition. SEA configuration and enablement is done on the VIOS partition. Implementing the SEA can help reduce the overall quantity of physical Ethernet adapters, cabling, and physical Ethernet switch ports required to support an environment. SEA requires VIOS to be configured with at least one physical Ethernet adapter port and one or more virtual Ethernet adapters.

4.7 HOST ETHERNET ADAPTERS (HEA)

A Host Ethernet Adapter (HEA) is a physical Ethernet adapter that is integrated directly into the GX+ bus on a managed system. The HEA provides lower latency in communications as well as higher efficiency, and therefore a potentially higher throughput. Unlike a standard Ethernet adapter, an HEA can't be assigned to a specific partition. Only Logical Host Ethernet Adapters (LHEAs) can be assigned to client LPARs. An LHEA is a representation of a physical HEA port on a logical partition. It appears to the OS as if it were a single-port physical Ethernet adapter.

The LHEA can be configured through the HMC GUI. In such a configuration, the virtual I/O server is completely unaware of the existence of an HEA on the system. This type of network virtualization provides some advantages over using a Shared Ethernet Adapter. It lowers the amount of resources used by the virtual I/O server and provides Ethernet connectivity with lower latency than traditional Ethernet adapters. There are two main drawbacks to using this type of configuration: the current inability to perform Live Partition Mobility of the client LPAR; and the lack of centralized management of link aggregation (if you use link aggregation).

4.8 N_PORT ID VIRTUALIZATION (NPIV)

NPIV is a Fibre Channel industry standard method for virtualizing a physical Fibre Channel port. NPIV allows one F_Port to be associated with multiple N_Port IDs, so a physical Fibre Channel HBA can be shared across multiple guest operating systems in a virtual environment. Each virtual FC HBA has a unique and persistent identity. On POWER® systems, NPIV allows logical partitions to have dedicated N_Port IDs, giving the operating system a unique identity to the storage system, just as if it had a dedicated physical HBA. The POWER Hypervisor generates WWPNs in pairs based on the range of names available for use with the prefix in the vital product data on the managed system. The hypervisor does not reuse the WWPNs that are assigned to the virtual Fibre Channel client (VFC) adapter on the client logical partition. The VIOS can support both NPIV and vSCSI simultaneously. Each physical NPIV-capable FC HBA supports 64 virtual ports.

When setting up an NPIV, consider the following points:

- The WWPN pair is generated each time you create a VFC and is never re-created or reused—just like a real HBA.
- Save the partition profile with VFCs in it. Make a copy; don't delete a profile with a VFC in it.
- Make sure that the partition profile is backed up for local and disaster recovery; otherwise it will be necessary to create new VFCs and map to them during recovery.
- For LPM to work, target storage subsystems must be zoned and visible from source and destination systems.
- For LPM to work, active-passive storage controllers must both be in the SAN zone.
- Do not include the VIOS physical adapter WWPNs in the zone.
- There should not be any NPIV LUNs in the VIOS.

- Load multipath code in the client LPAR, not in the VIOS.
- Monitor VIOS CPU and memory.

5 NETAPP STORAGE AND VIRTUALIZATION

5.1 AGGREGATE AND FLEXIBLE VOLUME (FLEXVOL)

NetApp aggregate and flexible volume (FlexVol[®]) technology delivers true storage virtualization solutions that can lower overhead and capital expenses, reduce disruption and risk, and provide the flexibility to adapt quickly and easily to the dynamic needs of the enterprise. Aggregate technology virtualizes storage resources through the creation of large pools of disks and enables creation of multiple flexible volumes on those aggregates. This flexibility means that customers can focus on managing data, not hardware, and can make changes quickly and seamlessly. As a result, customers can add storage when and where it is needed, without disruption and at the lowest incremental cost. For more information about creating and managing FlexVol volumes, refer to the `vol` command in the appropriate NetApp documentation, available online at NOW[™] (NetApp on the Web).

5.2 THIN PROVISIONING

Thin provisioning is a way of presenting more storage space to the hosts connecting to the storage system than has been physically allocated. With thin provisioning, the storage administrator takes an aggregate and creates volumes for different applications to use, but does not actually allocate any space upfront to those volumes. The space gets allocated only when the host application needs it. The unused aggregate space is available for the existing thinly provisioned volumes to expand or for use in creating new volumes. When the aggregate free space declines to a predefined threshold, the storage administrator can attempt to reclaim space from volumes that have obsolete data, can delete volumes that are no longer in use, or can add more storage disks. By allowing as-needed provisioning and space reclamation, thin provisioning can result in better storage utilization and smaller capital expenditures on the storage infrastructure.

5.3 FLEXCLONE

NetApp FlexClone technology instantly creates logical copies of data volumes and data sets as transparent, virtual copies. By implementing FlexClone, customers can increase productivity and save storage space without compromising performance. FlexClone technology can create true clones instantly without requiring additional storage space. Each cloned volume is a virtual copy that can be used for essential business operations such as creating development and testing environments, setting up environments to test the Fixpack and platform upgrades, and so on. FlexClone volumes are created in a matter of seconds and save space with minimal overhead, so you can afford to create all the clones you need to improve productivity.

6 HARDWARE CONFIGURATION

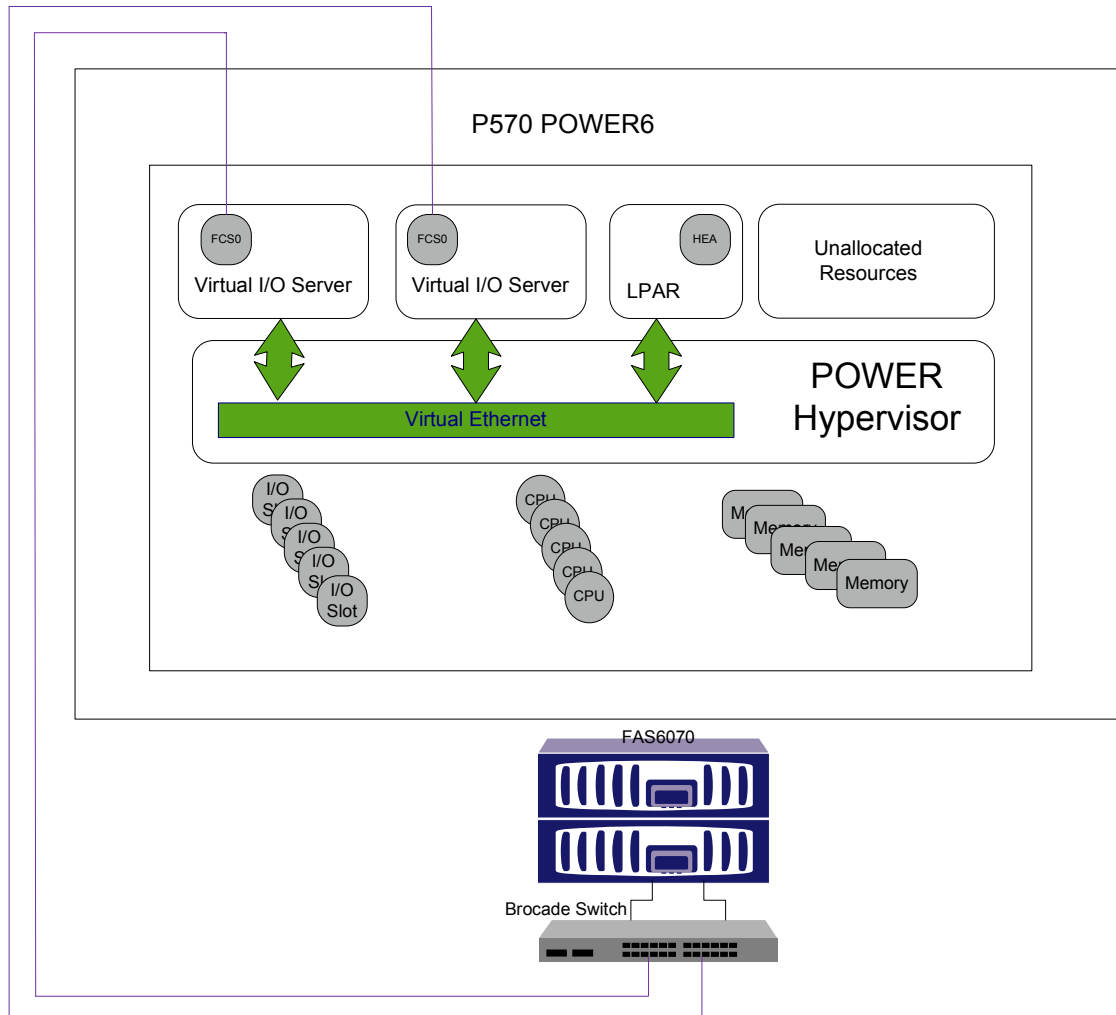


Figure1) Architecture of P570 and FAS6070C.

For this setup, an IBM P570 with two building blocks was used to set up the server infrastructure. The system includes the industry-leading IBM POWER6 technology and mainframe-inspired reliability, availability, and serviceability (RAS) features, as well as innovations like EnergyScale™ and PowerVM virtualization technologies. This innovative approach enables secure, nondisruptive growth, while helping to maintain outstanding performance and maximizing your investment.

A NetApp FAS6070C was used as the underlying storage system. The FAS6070 is a highly scalable storage system designed for the largest applications and consolidations in enterprise data centers. It is expandable to 1,008 disk drives and a maximum capacity of 504TB. The FAS6070 is available in either a single-controller configuration or an active-active, dual-controller configuration. Cost-saving SATA and performance FC disk drives can be used together in the same system for tiered storage.

For this setup, the FAS6070C was presented with three disk shelves on each storage controller. An 8Gb Brocade SAN switch with NPIV capability was used to connect the P570 and FAS6070C; this is a requirement for the virtualization of Fibre Channel adapters from the virtual I/O server to the AIX LPAR. Figure1 shows the P570 with FAS6070C that was used to set up this environment. For detailed lists of configuration components, see the appendix.

7 ENVIRONMENT CONFIGURATION

7.1 VIRTUAL I/O SERVER

The Virtual I/O Server software resides in a logical partition and requires its own resources, including processor, memory, and dedicated disk. The software is installed on a dedicated disk and is called the Virtual I/O Server's *root volume group*. Because the VIOS owns hardware adapters like SCSI disks, Ethernet, Fibre Channel adapters, and optical devices, it is essential that it always be available. To make sure that the client LPARs gets all the required resources from the VIOS, NetApp recommends:

- Create two VIOS servers to serve the client LPARs.
- When booting the Virtual I/O Server from locally attached SCSI disks, the root volume group should extend to include a second SCSI disk and the data mirrored on that disk.
- Each Virtual I/O Server is presented with an 8 GB physical Fibre Channel adapter.
- As shown in Figure 2, four virtual Fibre Channel server adapters are created for each VIO client.
- Two virtual Fibre Channel server adapters are mapped to each physical Fibre Channel adapter.

Benefits of these recommendations include the following:

- Having more than one VIOS creates redundancy, which allows easier system management, such as software upgrades. With two Virtual I/O Servers, the layer of physical resource dependency has been removed.
- Mirroring the root volume group of each VIOS to multiple disks creates redundancy at the disk level, which means high availability of each individual VIOS.
- Mapping multiple physical Fibre Channel adapters and creating multiple virtual Fibre Channel adapters on each physical adapter creates adapter-level redundancy.




Figure2) Virtual Fibre Channel setup through NPIV.

7.2 NETWORK CONNECTIVITY

To access the VIO client externally requires network connectivity, which can be configured by using either an HEA or an SEA. For this setup we used multiple HEAs to provide network redundancy, as shown in Figure 3. NetApp recommends the following for this configuration:

- Two LHEA ports assigned to the VIO client
- A link aggregation on the client LPAR using the two LHEA ports

Benefits from the above recommendation for network configuration are:

Using LHEAs simplifies Ethernet connectivity in a couple of ways. It makes it possible to avoid the more complex VIOS SEA configuration and to avoid VIOS dependency for network connectivity. Because an LHEA is virtualized at the hypervisor level, it provides lower latency in communications as well as higher efficiency, and therefore provides potentially higher throughput.

Using more than one physical LHEA and creating the link aggregation achieves the desired redundancy. Each LHEA physical port can have up to 16 logical ports, providing sufficient Ethernet ports.

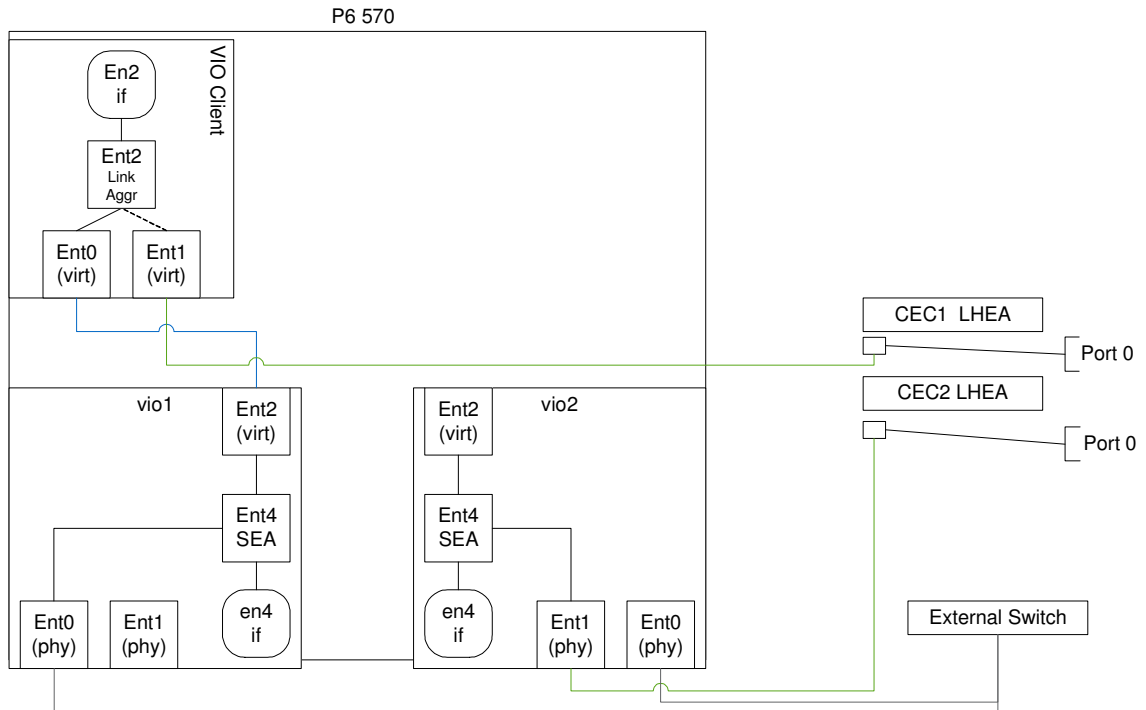


Figure3) Ethernet setup for the Virtual I/O Server and Virtual I/O client

7.3 AIX LPAR (VIO CLIENT)

For our test environment, we created one VIO client LPAR using six CPU units and 4GB of RAM. We then installed AIX 6.1. The VIO client is served by dual VIOS. The purpose of the VIOS is to serve the VIO client with the required Fibre Channel adapters. For storage we created a total of 10 LUNS on the NetApp storage system, sharing the same NetApp aggregate. The LUNS were mapped from the NetApp storage system through the VFC adapters of the VIOS and the VIO client. The JFS2 file systems was created on each AIX volume group, with each being mapped to its own NetApp LUN. All other resources required by the VIO client are served at the hypervisor level, resulting in less dependency on the VIOS. We took the following steps for maximum availability of required resources to the VIO client:

- The VIO client LPAR was SAN booted through a NetApp LUN.

- Four virtual Fibre Channel client adapters were created and presented by two physical adapters, one from each VIO Server. Figure 2 illustrates the VFC adapter setup.
- Each LUN was mapped to all four virtual Fibre Channel client adapters.
- We used AIX native MPIO to achieve multipathing.

7.4 AIX TUNING PARAMETERS

We used the following nondefault AIX parameter settings to maximize throughput during SIO workload tests on the VIO client:

- **minperm%.** The minimum amount of memory the persistent file cache can occupy was set to 5% and was modified by the following command:

```
vmo -o minperm%=5
```

- **maxperm%.** The maximum amount of RAM the persistent file cache can occupy before it is used as the sole source of new pages by the page-stealing algorithm was set to 90%, and was modified by the following command:

```
vmo -o maxperm%=90
```

- **queue depth.** The maximum number of requests the disk device can hold in its queue. NetApp recommends setting the queue depth on each hdisk to 256 by the following command:

```
chdev -l <device name> -a queue_depth=<value>
```

For example, to set the queue depth for hdisk0 to 256, you use the following command:

```
chdev -l hdisk0 -a queue_depth=256
```

- **max_transfer.** The maximum size of a transaction in bytes that the disk driver will issue to the adapter driver. It can be modified by the following command:

```
chdev -l hdisk1 -a max_transfer=<value>
```

NetApp recommends setting the `max_transfer` value to 1MB. To set this value for hdisk1, use the following command syntax:

```
chdev -l hdisk1 -a max_transfer=0x100000
```

These changes can be confirmed by issuing the `lsattr` command. To verify the changes for the hdisk named hdisk0, use the following command:

```
lsattr -El hdisk0
```

The benefits from these recommendations are:

- Installing AIX on a NetApp SAN LUN provides access to NetApp data protection, disaster recovery, and dev/test technologies, such as Snapshot™ and FlexClone, for high availability.
- Modifying the AIX parameters (`maxperm` and `minperm`) means better use of system memory during file system caching.
- The `queue depth` and `max_transfer` settings support optimum data transfer. These parameters were changed based on the workloads.

7.5 FIBRE CHANNEL ADAPTERS

The following `queue depth` settings were applied on the physical and virtual Fibre Channel adapters. The `queue depth` parameter controls how much data is allowed to be “in flight” on storage from the HBA. If you have I/O-dependent applications that use SQL, NetApp recommends changing the default setting to allow more in-flight I/O requests.

On both VIOSs, the queue depth is set to 2048 for all physical Fibre Channel adapters by the following command:

```
chdev -l <adapter name> -a num_cmd_elems=2048
```

To set the queue depth to 2048 for the physical adapter named fcs0, use the following command:

```
chdev -l fcs0 -a num_cmd_elems=2048
```

On the VIO client LPARs, the queue depth for each virtual Fibre Channel adapter was set to 1024. To set the queue depth to 1024 for the virtual Fibre Channel adapter named fcs2, use the following command;

```
chdev -l fcs0 -a num_cmd_elems=1024
```

The queue depth changes can be confirmed by issuing the `lsattr` command. To verify the queue depth changes for the adapter named fcs0, use the following command:

```
lsattr -El fcs0
```

To summarize, implementation of these queue depth recommendations on both the VIO Server and client optimizes data transfer.

7.6 SAN SWITCH

For this setup, an 8Gb Brocade switch was used and the N_Port ID Virtualization feature (NPIV) was enabled to allow virtualization for assigned ports. Instead of using auto-negotiate, port speeds were set to 8Gb and 4Gb for the AIX LPAR and the storage system respectively.

ENABLING NPIV

To enable NPIV capability on each port for the SAN switch, use the `portCfgNPIVPort` command. For example, to enable NPIV on port 2, use the following command:

```
labsan01:admin> portCfgNPIVPort 2, 1
```

The NPIV capability on the SAN switch can be verified by using the `portcfgshow` command:

```
labsan01:admin> portcfgshow
```

The next step is to do the appropriate zoning on the SAN switch by using the WWPN of the virtual Fibre Channel client adapter created in the virtual I/O client partition. To do the zoning with the correct WWPNs, refer to your SAN switch manual.

7.7 NETAPP STORAGE SYSTEM

The following configuration was used for this setup:

- One large aggregate was created for each storage controller. Each aggregate was created with two RAID groups, each with 19 disks, for a total of 38 disks. This configuration was chosen to get the maximum throughput based on the number of spindles.
- One FlexVol volume was created for booting the AIX LPAR and was named `aix_boot_vol`.
- One FlexVol volume was created for data and was named `aix_test_data`.
- Flexible volumes (FlexVol volumes) provide optimal I/O performance by automatically striping data across all disks in an aggregate.
- A RAID-DP[®] configuration was used for this test environment. NetApp storage systems support two RAID configurations, RAID-DP and RAID 4. For technical details, read the white paper [RAID-DP: NetApp Implementation of RAID Double Parity for Data Protection](#).
- The following options were modified for each flexible volume:
 - **Disabled automatic NetApp Snapshot copies.** By default, when a volume is created, a default Snapshot schedule is set for it. A backup schedule for this particular environment was not required. However, if a backup schedule is required, you must back up the database based on a user-defined schedule. Therefore automatic Snapshot scheduling was disabled for the tests on each volume by running the following command on the storage systems:

```
vol options [VolName] nosnap on
```

Where:

`VolName` is the name of the volume.

- **Disabled access time update of files.** For volumes that are used by databases, the database management system (DBMS) manages the correct file-access time for the inode. Therefore access time update by storage is not required, and this option was enabled by running the following command for each volume:

```
vol options [VolName] no_atime_update on
```

Where:

`VolName` is the name of the volume.

- **Set the nvfail option on.** If this option is on, the NetApp storage system performs additional status checking at boot time to verify that the storage system's nonvolatile RAM (NVRAM) is in a valid state. This option is useful for databases because any problems found with NVRAM result in the database instances failing to start and error message being sent to the console to alert database administrators (DBAs). The following command was used to set the `nvfail` option for each volume:

```
vol options [VolName] nvfail on
```

Where:

`VolName` is the name of the volume.

- **Set the Snapshot reserve.** When a new volume is created, by default Data ONTAP® reserves 20% of the space for Snapshot copies that cannot be used for data. NetApp recommends setting the Snapshot reserve based on your actual Snapshot space requirement, which depends on the backup retention policy. The Snapshot space requirement can be computed by using the formula (Actual space consumed * Age of the oldest Snapshot * rate change). To better use the storage space, for this test environment the Snapshot reserve was set to 0 on each volume by running the following command:

```
snap reserve -V [VolName] 0
```

Where:

`VolName` is the name of the volume.

- **Set snapshot autodelete.** Snapshot autodelete is a policy-based space-management feature that automatically deletes the oldest Snapshot copies on a volume when that volume is nearly full:

```
snap autodelete [VolName] on
```

- **Set the fractional reserve.** When a new volume is created, by default Data ONTAP sets the fractional reserve to 100%. If you have enabled space reservation for a LUN, you can reduce the space that you preallocate for those reservations by using fractional reserve. NetApp recommends setting the fractional reserve to 0 on each volume by running the following command with snapshot autodelete turned on:

```
vol options [VolName] fractional_reserve 0
```

Where:

`VolName` is the name of the volume.

In addition to these storage settings, you can use the NetApp storage virtualization features such as thin provisioning and FlexClone to save storage space and realize cost savings.

8 WORKLOAD TEST

To validate the setup an SIO workload test was conducted with different workload types. SIO is a tool from NetApp that generates artificial I/O workloads against any device. It supports numerous configuration variables (reads versus writes and so on). It supports multiple devices and multiple threads, and it collects a wide variety of statistics on I/O client machines and/or I/O servers.

To increase I/O throughput, multiple files were created for SIO workload testing. A total of 30 files were created, 3 under each file system. The SIO workload test was run using the `-direct` option to avoid JFS2 caching. To exploit storage system caching, the total file size during the read operation was kept under 32GB, which is the size of the storage system memory.

Figure 4 shows the SIO workload test done with 100% sequential reads and varying thread counts. The test shows a throughput level of 820MB/s. This number almost matches with the Performance Quick Report (PQR) report for a single-head FAS6070 storage system.

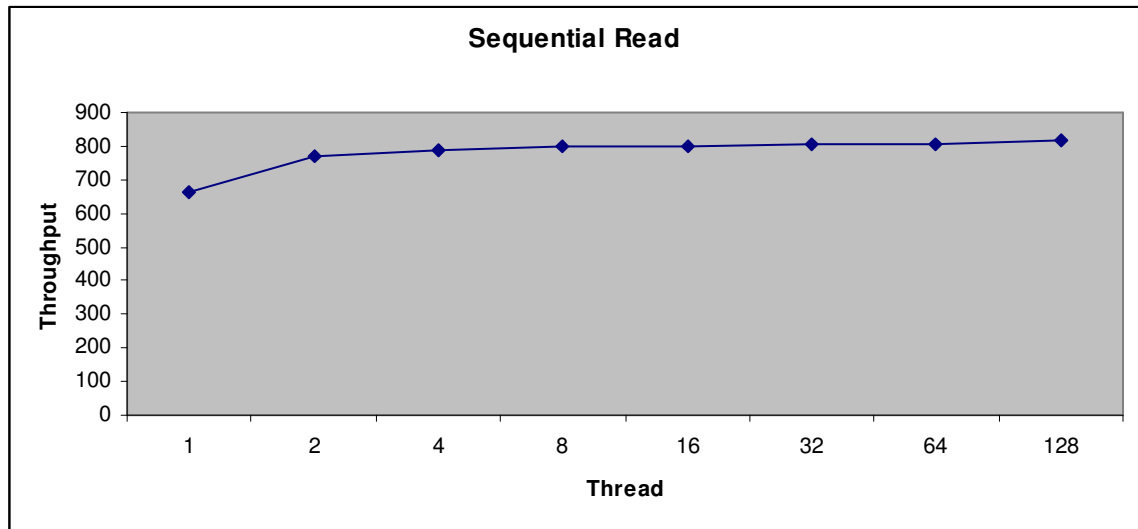


Figure 4) Throughput measurement on 100% sequential reads.

Figure 5 shows SIO workload test results with 100% random reads varying thread counts. The test shows a maximum throughput level of 292MB/s.




Figure 5) Throughput measurement of 100% random reads.

Figure 6 shows the results from an SIO workload test done with 100% sequential writes and varying thread counts. The test shows a maximum throughput level of 429MB/s.

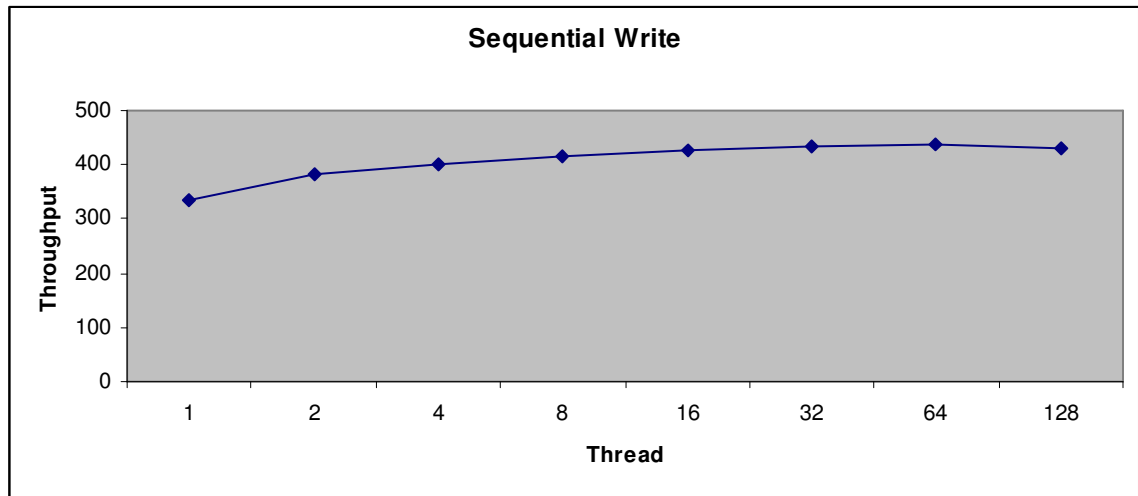


Figure 6) Throughput for 100% sequential writes.

Figure 7 shows the results from the SIO workload test done with 100% random writes and varying thread counts. The test shows a maximum throughput level of 145MB/s.

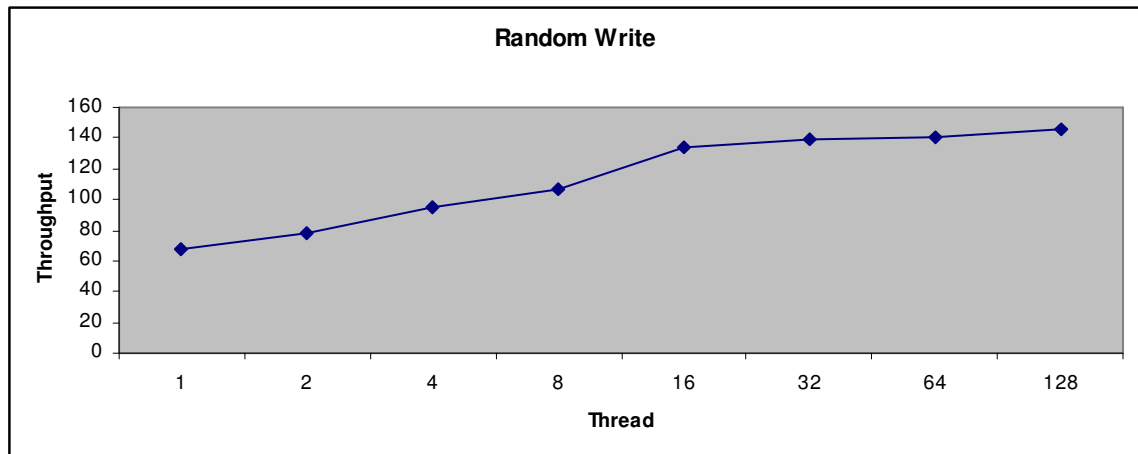


Figure 7) Throughput for 100% random writes.

9 CONCLUSION

The workload tests done on the AIX LPAR with NetApp as the underlying storage system gave a throughput that almost matches the PQR report of FAS6070. The workload test also proves that the overhead of virtualizing the FC adapter is minimal. The IBM PowerVM, together with the NetApp storage system, offers customers a compelling advantage in terms of virtualization. Using the PowerVM technology, the customer can do the virtualization of the server resources such as CPU, memory, Fibre Channel devices, network adapters, and so on. Using NetApp as the underlying storage system, the customer can exploit the storage level virtualization by using NetApp technologies such as FlexVol, FlexClone, and thin provisioning.

10 ACKNOWLEDGEMENT

Igor Feoktistov, UNIX System Administrator, Engineering and Standards, NetApp.

Dennis P. Murphy, Senior System Engineer, Tech and Integration Delivery, NetApp
John Elliot, Database Performance Engineer, NetApp

APPENDIX

Power6 570 system details

Component	Details
<ul style="list-style-type: none">Processors	<ul style="list-style-type: none">16
<ul style="list-style-type: none">Total physical RAM	<ul style="list-style-type: none">64GB
<ul style="list-style-type: none">Ethernet connection	<ul style="list-style-type: none">Two 4-port 1Gb Host Ethernet Adapters
<ul style="list-style-type: none">Fibre Channel connection	<ul style="list-style-type: none">Six 8Gb PCI Express Dual Port FC Adapters
<ul style="list-style-type: none">Firmware	<ul style="list-style-type: none">EM340_041

VIOS Details

Component	Details
<ul style="list-style-type: none">Processor	<ul style="list-style-type: none">Two shared processing units (uncapped)Four virtual processors
<ul style="list-style-type: none">Total physical RAM	<ul style="list-style-type: none">2GB
<ul style="list-style-type: none">Fibre Channel connection	<ul style="list-style-type: none">Four 8Gb Virtual FC Adapters
<ul style="list-style-type: none">Version	<ul style="list-style-type: none">2.1.1.10-FP21 + APAR IZ56171

AIX LPAR Details

Component	Details
<ul style="list-style-type: none">Processor	<ul style="list-style-type: none">Six shared processing units (uncapped)Eight virtual processors
<ul style="list-style-type: none">Total physical RAM	<ul style="list-style-type: none">4GB
<ul style="list-style-type: none">Fibre Channel connection	<ul style="list-style-type: none">Four 8Gb Virtual FC Adapters
<ul style="list-style-type: none">Version	<ul style="list-style-type: none">AIX 6.1 TL3 SP1 + APAR IZ56171

NetApp FAS6070 Details

Component	Details
<ul style="list-style-type: none">Processor	<ul style="list-style-type: none">Four 2.6GHz
<ul style="list-style-type: none">Total physical RAM	<ul style="list-style-type: none">32GB
<ul style="list-style-type: none">Fibre Channel connection	<ul style="list-style-type: none">One dual port 4Gb FC adapter for the host

	<ul style="list-style-type: none"> • One quad port 4GB FC adapter for the back-end disk shelves
<ul style="list-style-type: none"> • Version 	<ul style="list-style-type: none"> • Data ONTAP 7.3.1

Brocade Switch Details

NetApp storage system details.

Component	Details
<ul style="list-style-type: none"> • Speed 	<ul style="list-style-type: none"> • 8GB
<ul style="list-style-type: none"> • NPIV 	<ul style="list-style-type: none"> • Capable
<ul style="list-style-type: none"> • Fabric OS version 	<ul style="list-style-type: none"> • v6.1.1b

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein must be used solely in connection with the NetApp products discussed in this document.