Do virtualization and SSD prove advantageous for Storage products with high speed interfaces in a SAN Environment?

Dr. M. K. Jibbe

Distinguished Engineer

Manager and Technical Lead of Test Architect and Product Certification in India

LSI Corporation (Engenio Storage Group)

Agenda

- Storage Area Network (SAN) – Introduction
- Storage Virtualization
- Efficient design of virtualization as it applies to Storage products
- Solid State Disks (SSD) in SAN
- Integration of SSD technology in a Storage Array Architecture
- SSD & SAN Virtualization
- Optimizing SAN using SSD in Virtualized Environment
- Optimization areas in Virtualization to fit a storage architecture
- Handling of SSD in a group of other disk drives
- High Speed 8Gb FC in Virtualized Environment
- Future of Virtualization: Predictions
- Summary/Conclusion
- References
- Backup Slides

Storage Area Network (SAN)

- Deployed as a powerful means of controlling the escalating cost and complexity of data
  - Administration
  - Management
  - Movement
- Initial Primary Focus on Issues like
  - Server-less backup
  - Storage consolidation
  - Ease of management
  - High availability operation
- Impressive application performance gains can be realized by incorporating Solid-State Disks (SSDs) into SANs
- Maximum results are achievable in a virtualized SAN environment with low power, cost and management.

SAN : An Efficient Use of Resources

- A typical enterprise storage plan should predict application failures due to insufficient space and consider
  - projecting storage capacity
  - building in extra capacity
  - Application predictability
  - Risk the Storage Administrator chooses to take
- This surplus capacity typically ranges from 10% to a whopping 50% of total storage capacity
- Compelling reasons to Implement SAN
  - Administrative personnel
  - More efficient use of Storage resources
  - Mitigate the headaches of DAS
  - No dealing with Separate storage per server
- One-to-many configuration allows for
  - More efficient capacity provisioning
  - Economical use of storage administrators' time
  - Improved reliability, availability, and scalability.

Storage Virtualization

- Even though basic storage devices may be getting cheaper, complexity is not just frustrating; it's also expensive.

- The solution is Storage Virtualization.
  - Simplifies Storage provisioning
  - Reduces administrative overhead
  - Enables and simplifies the targeted provisioning of resources

- What is storage virtualization?
  - Layer of abstraction that separates the representation of storage from the physical devices themselves
  - Provides the ability to consolidate multiple physical storage devices of various interface protocols (such as SCSI, iSCSI, Fibre Channel) into a logical "storage pool."

- Virtual storage devices can be created from Storage Pool and benefits are
  - Easier storage resource management for the system operator
  - Maximized utilization of all physical storage space
  - Improved storage network performance
  - Elimination of downtime to add and reconfiguration storage resources
  - Automated capacity-on-demand for assigning storage based on policy-driven strategies

Getting the Most Out of SAN Storage Virtualization

- Storage capacity projection is radically reduced compared to non-virtualized SANs
  - Idle or unused storage is eliminated as a result of the pooling of *all* the available storage capacity
  - Administrators can be proactive, rather than reactive, regarding necessary changes to storage.
  - Policies add storage capacity automatically when pre-defined thresholds are reached.
- In non-virtualized storage farms
  - Physical co-location of storage devices
  - Each device is managed, configured, and licensed separately
  - Results in vendor incompatibilities in management capabilities.
- In comparison, in-band storage
  - Virtualization offloads virtualization and other storage services from the servers and storage devices
  - Puts the intelligence into the network.
  - Storage Virtualization allows enterprises to consolidate all their storage (JBODs, RAID arrays, tape drives/libraries, SSDs etc.)… into a *unified pool of heterogeneous storage*, independent of
    - hardware type
    - vendor
    - protocols
- The result is a centralized storage solution including
  - Services
  - Configuration
  - Management that delivers accelerated performance and boundless scalability.

Storage Virtualization Demonstration

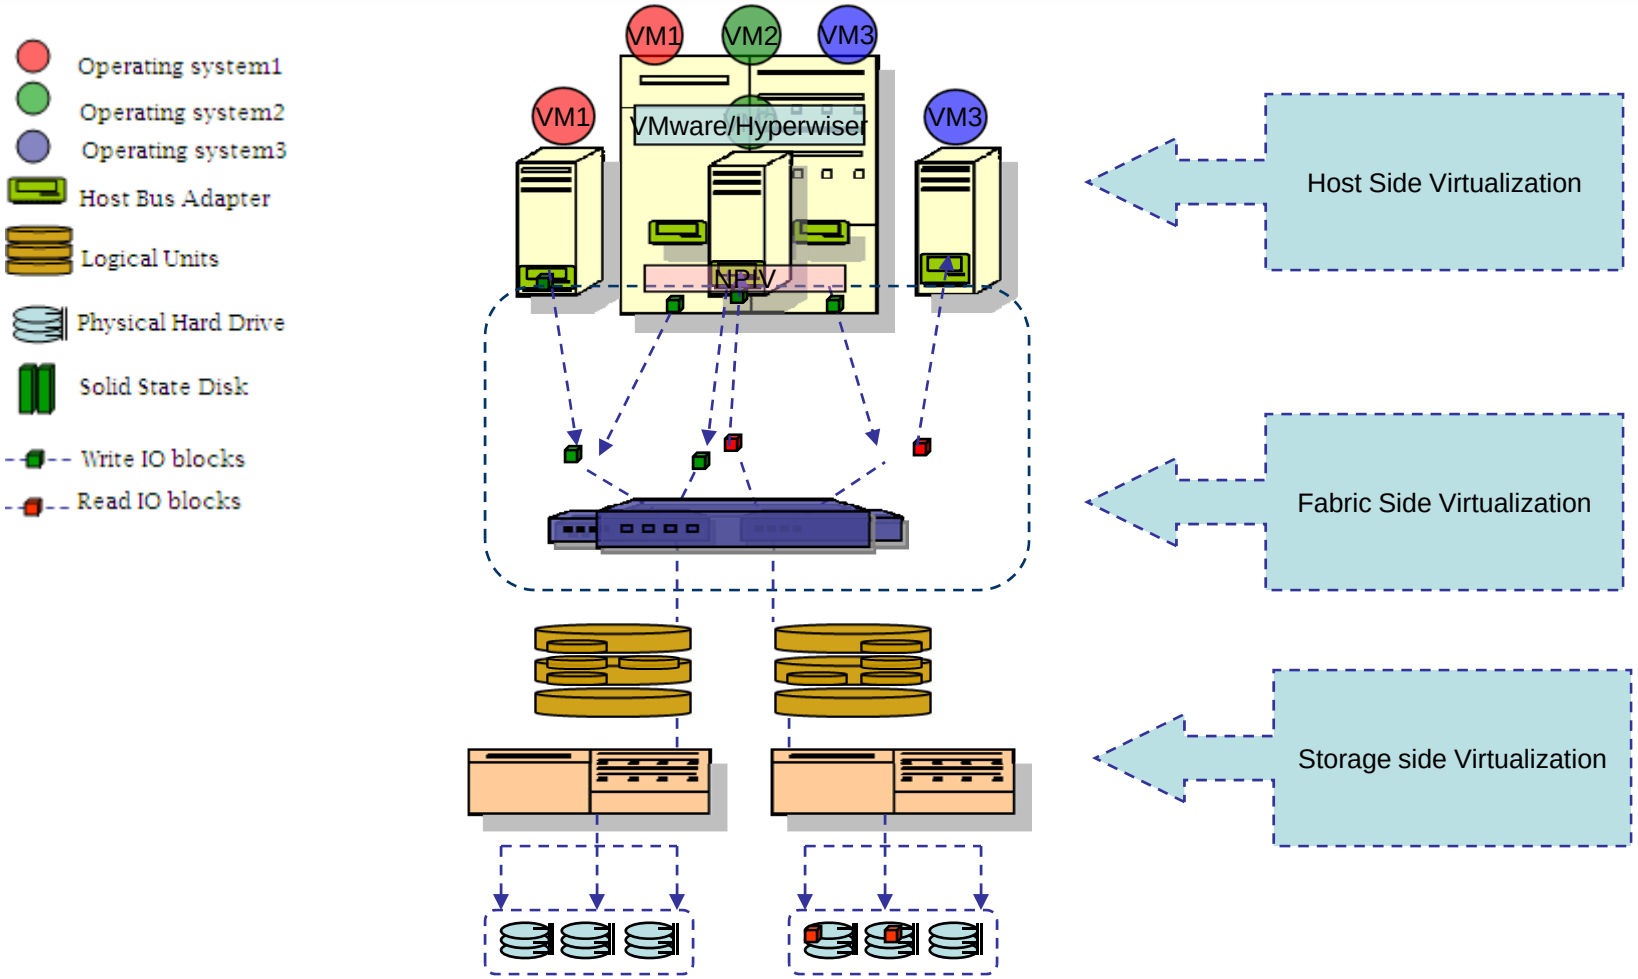

SAN-based storage virtualization

- Virtualization is what modern storage systems are all about!
- Host-based volume managers
  - Storage is connected to external cached disk arrays
  - hides the messy details of SCSI targets, controllers and RAID protection from the servers.
- Unresolved Storage Management Problems
  - Tough to migrate data from one storage platform to the next
  - Administrative burden of storage systems is still too manual and difficult
  - Most of the time utilized to keep the storage running instead of Optimizing it
- SAN-based storage virtualization holds the promise of addressing these problems, enabling you to improve your customer's data management and deliver storage efficiencies.
- Empowering customers to migrate and manage data quickly means that they will be more inclined to refresh aging technology.
- This tip covers five value propositions for SAN virtualization.
  1. Single point of administration
  2. Non disruptive data migration
  3. Information lifecycle management (ILM)
  4. Improved allocation efficiencies
  5. Heterogeneous replication

Efficient Design - Mistakes to avoid in Virtualization

- Capacity and costs fail to meet expectations
- **Server failure.** Large-scale consolidation may put many key processes, applications and services in the same proverbial basket. Consequently, fewer physical servers bear the workload -- and a physical failure has much more significant consequences.
- **Over-provisioning**. Starting consolidation without a clear picture of an application's function, workload or profile may lead to infrastructures that are out of balance and over-provisioned.
- **Service levels.** Virtualization technology requires new skills; for example, the ability to identify whether a problem originates in the physical or virtual environment. Without staff skills to address problems, service levels may suffer.
- Hardware bandwidth to accommodate the increased virtual load
- Platform to support allocating resources for each virtual load
- Hardware downtime should not cause downtime for entire virtual platforms on top of it. The platform should support hot insert and replace on resources allocated to mitigate the downtime for all virtual load
- Backup and disaster recovery implementation gets more complex and rapid recovery becomes more important in a virtualized environment

8Gb FC in virtualized environments

- The modern data center is faced with ever-growing demands
- 8Gb FC HBAs and Fabric offer a multitude of hardware-assist virtualization enhancements that help provide a complete HBA and Fabric virtualization solution
  - Simplify management and provide port-level QoS
  - Reduce system overhead and improve reliability
  - Achieve higher, deterministic I/O performance with QoS functionality
    - Multiple device I/O queues
    - Unique requester IDs
    - Improved device interrupt routing.

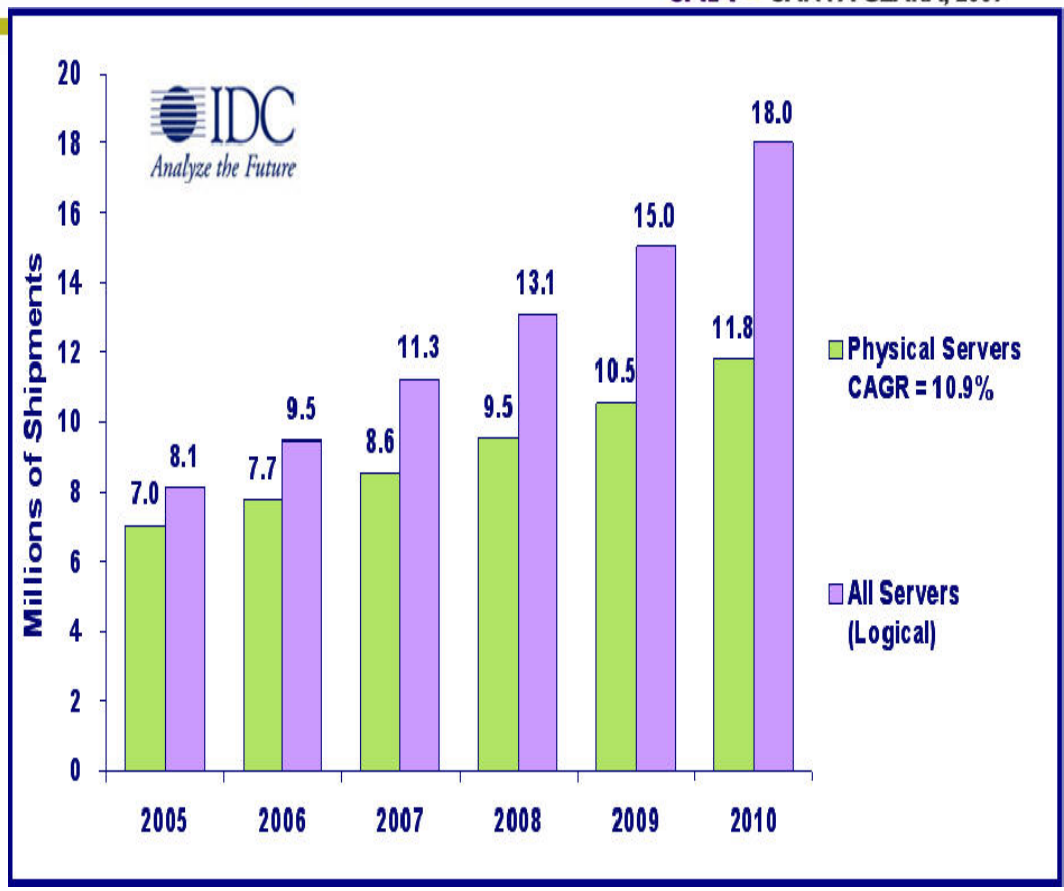

- Less than 5% of servers today are virtualized
- CAGR of close to 157% from 2005-2010.
- Industry estimates: >70% of the virtualization deployments are attached to Fibre Channel SANs.

8GFC HBAs in virtualized environments

- 8GFC NPIV implementation virtualizes the physical HBAport configured in a point-to-point SAN topology
- A single physical FC port functions as multiple logical ports, each with its own identity.
- The benefits of Virtual ports to storage administrators include:
  - Binding VMs to storage and defining multiple zones
  - VMs access restriction is independent of the physical port access
  - Eliminating the need to reconfigure zoning and LUN masking
  - Enabling storage administrators to continue to use their typical SAN management applications to provide
  - Reducing the overhead associated with software hypervisor by offloading specific processing into the hardware.

Solid State Disks (SSD) in SAN

- Storage devices most commonly deployed in SAN implementations include
  - Cached disk arrays
  - JBOD
  - Tape drives
- A key SAN benefit is the ability to make more efficient use of any and every kind of storage resource.
- SSDs or file cache accelerators can be used as a Storage resource
- SSDs have historically been used to successfully respond to
  - I/O bottlenecks affecting application performance on a single server
  - Single application basis
- Typically, 3% to 5% of application data accounts for approximately 50% of all I/O activity.
- Isolating these "hot files" on an SSD typically improves
  - Individual file-level performance 200-800% and
  - Commonly increases overall application performance by 30-40%.

SSD in the SAN

- SSD as a dedicated device for the broader spectrum of applications has been historically difficult

- Leading SSD devices
  - Support up to 8 or more Fibre Channel ports for switched SAN environments
  - Partitionable into numerous logical unit numbers (LUNs)
  - Can be managed using standard off-the-shelf SAN and SRM tools.

- These capabilities allow the performance benefits of SSD to be effectively
  - Leveraged across the SAN and specifically directed to areas of a SAN where performance acceleration is most needed, when it's needed
  - Enables SSDs to be increasingly used to supercharge applications when only periodic high speed disk is desired for peak times
  - Used in more traditional applications that need full time performance enhancement, such as database logs or message queuing systems..

SSD Pros & Cons

Pros

- No initial spin-up required - Significantly faster startup
- Even the fastest Hard drives today do not have half the read time of the SSD
- SSDs have seek speeds hundreds of times faster than that of mechanical disks
- Since there are no moving parts in SSDs, there is literately no sound emitted from the device
- Greater reliability
- Energy efficient. "Using SSD technology reduces the overall power consumption of devices such as disk arrays and also improves their performance and environmental ruggedness
- Durability. "SSD is designed to operate in more extreme environments of up to 70 degrees Celsius

Cons

- On a price per gigabyte basis, SSDs are considerably higher than conventional hard disks;
- Life expectancy and Technology is nascent

SSD In Array Architecture

- SSD designs in an array environment
  - Cache on the host side
  - Write and Read ahead data cache in the array controller firmware
  - Meta data for an array controller to improve
    - Controller Start of Day
    - Exception handling
    - Error recovery timing

- Managing an SSD behind an array controller
  - Array controller designed to manage hard drive with SSD
  - Array controller firmware architecture tuned to utilize the performance to cost benefit of using SSD

SSD & Virtualization

- The coupling of virtualization and SSD in a storage array Architecture
  - Tight coupling between SSD and Array design to identify hot-spots and leverage performance benefits from SSD
  - SAN "Tuning" to address SAN QoS requirements
  - Power Savings in SAN environment using SSD

Optimizing SANs by Using SSD in a Virtualized Environment

- Storage virtualization eases the Administrator's burden
- Grants the ability to assign storage resources based on
  - Performance characteristics
  - application performance requirements
- This empowering capability opens up valuable possibilities for "SAN tuning"
  - To address Quality of Service (QoS) issues within the SAN.
  - Since the SAN infrastructure shares common components a valid concern is ensuring that high priority data requests are responded to promptly and not queued behind pending low-priority data requests (i.e. Quality of Service).
- Virtualization facilitates storage model within the SAN that responds to these QoS issues.
  - Seamlessly migrated between
    - Platforms
    - Operating systems
    - Applications for true point-and-click performance resolution.
- SSDs can be deployed into the virtualized SAN to dynamically deliver top performance and a statistically consistent level of data accessibility
- The virtualized SSD becomes an architectural resource to address performance issues

SSD in SAN Virtualization - Optimization

- For example,
  - The SSD can be utilized on January 31st to speed up month-end closing on an Windows based financial application.
  - On February 1st, the SSD (or a portion thereof) can be allocated to help the engineering department with a Unix-based OS conversion dynamically by the virtualization management capabilities.
  - This kind of mobility and scalability
    - Allows the organization to leverage the performance benefits of solid-state disk across a broader application spectrum
    - Recognize the benefits in virtually every facet of the enterprise
    - Reducing the total cost of ownership.
- The ultimate goal for the administrator is to not have to worry about SAN tuning at all. The virtualization engine
  - Monitors all data requests to the storage devices
  - Gathers statistics on the location of the data hot spots.
- Eventually, the ability to dynamically act on this information and move the most requested data to the fastest devices will provide automatic SAN tuning.
- This automatic SAN tuning ensures that the entire spectrum of applications will have access to the right storage at the right time.

SSD in SAN Virtualization - Demonstration

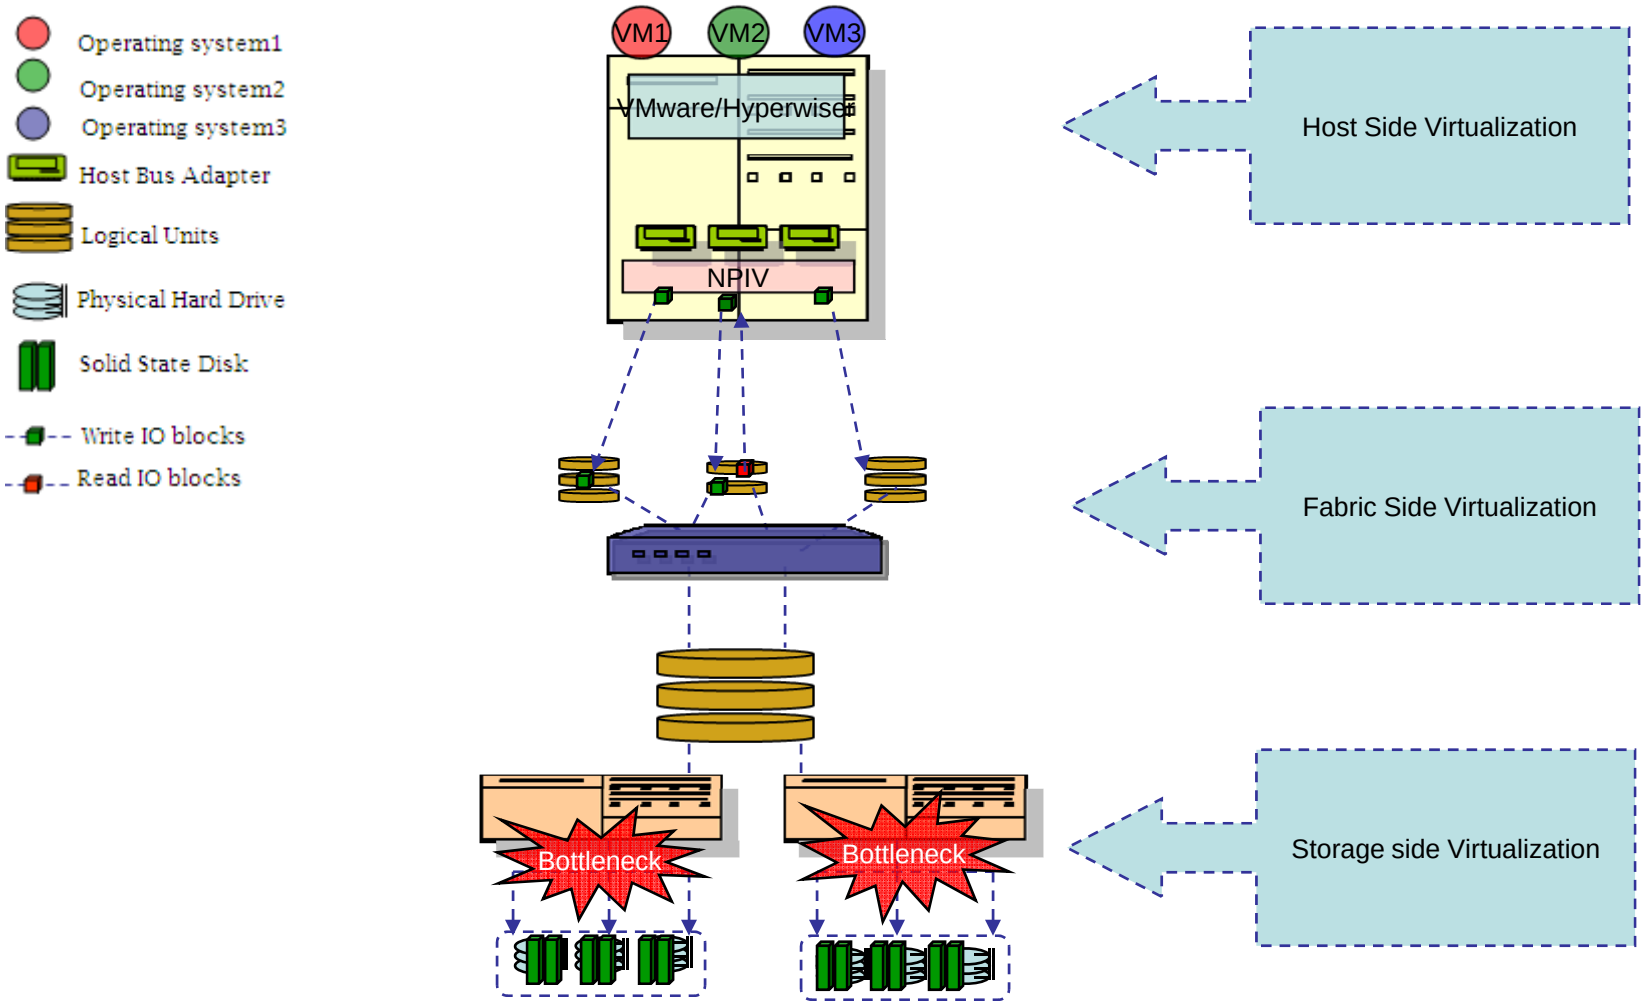

Optimization Areas in Virtualization

- Optimization areas in Virtualization to fit a storage architecture
  - Improve Storage capacity and resource utilization
  - Simplify management of Storage resources
  - Dynamic Storage allocation
  - Uninterrupted online migrations for Storage administrations

Handling of SSD in a group of other disk drives

- Enhanced Array architecture to use SSD with existing Hard Drive
- Developing Array firmware to be aware on Storage hot-spots and utilize performance benefits from SSD
- Simplified management for SSD and other disk drive usage
- Customized design to over come bottlenecks over Array virtualization

Fabric Access Challenges in VM Deployments

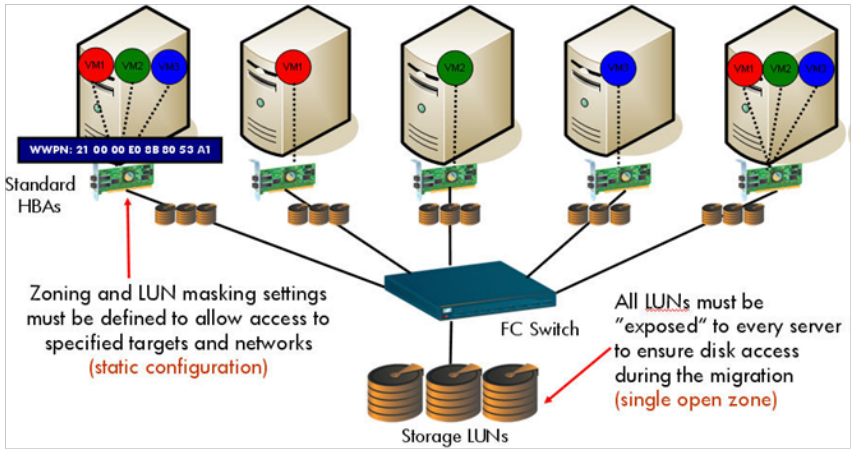

Fabric Access Solutions for VM Environments

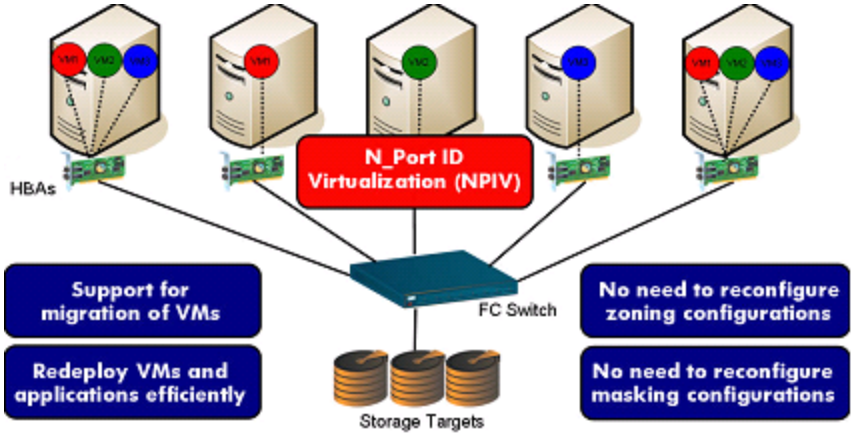

- SSD Storage enables a data center to take advantage of the high speed FC
- Designing solid-state disk into the virtualized 8 GFC SAN provides a powerful and cost-effective resource that can be used to significantly improve performance, productivity, and profitability across the enterprise.

Future of Virtualization

- FCoE
- SSD on virtual SAN
- Data Deduplication

Summary/Conclusions

- SANs are an enabling technology that is in the late adoption phase of the technology lifecycle.
- The networking aspects of SANs lend themselves to deploying other enabling technologies in an increasingly more powerful and broad-based role across the organization.
- With innovative storage virtualization tools that continue to evolve to better help manage storage environments, companies can focus on tuning the SAN for maximum benefit to all applications.
- Designing solid-state disk into the virtualized SAN provides a powerful and cost-effective resource that can be used to significantly improve performance, productivity, and profitability across the enterprise.
- Storage virtualization in all of its forms has been providing benefits to businesses for years; yet there are still some challenges that remain.
- It's easy to see that SAN-based virtualization holds a lot of promise for the industry.
- Discuss these value propositions with your customers to open their minds to the possibilities of SAN virtualization

Backup Slides

SAN Virtualization - Value Propositions

1. Single point of administration

- Customers understand that a little friendly competition between storage vendors can help reduce the price of storage
- SAN virtualization can be implemented in multi-vendor storage environments, so if a potential customer doesn't run the brand of storage equipment that you offer, converting them to a SAN virtualization setup could open the door to storage hardware sales.
- SAN-based storage virtualization provides this benefit as the administration occurs at the virtualization layer.

2. Non disruptive data migration

- Many customers do not replace their obsolete storage arrays
- It is very difficult to migrate to the next storage platform.
- Without storage virtualization, migrations often require application outages and lots of sweat equity.
- Coordinating effort between various teams/BU are difficult
- With SAN virtualization, the storage team can execute disk array swap-outs without impact
- Allows VARs to Sell new storage

SAN Virtualization - Value Propositions

3. Information lifecycle management (ILM)

- Customer wants the data to be in appropriate tier of storage
- Data access patterns are a key criterion in determining where to put the data
- Frequently accessed and require high-performance storage
- Rarely accessed could exist on more cost-effective storage.
- SAN virtualization tools sit between the server and the storage hardware, they have awareness of the access patterns.
- Can relocate the data to appropriate storage type based on access
  - Frequently Accessed - more expensive, high-performance storage
  - Less Frequently Accessed - less expensive storage
- This brings true ILM within reach

4. Improved allocation efficiencies

- Improving asset utilization is a quick way to lower the total cost of ownership (TCO)
- Low utilization - demanding more storage than needed.
- Team/Project could be new and no history to Plan for data growth
- Storage virtualization promises to solve both problems.
  - Thin provisioning services in the virtualization layer
  - Allows pre-allocation of storage
  - Shared free space across applications optimizing unused disk

5. Heterogeneous replication

- Disaster recovery replication – Huge challenges associated with maintaining agnosticism among disk array vendors
- Most array-based storage replication is not heterogeneous,
- Host-based replication options are heterogeneous, but management is cumbersome when a large number of hosts have replicated data
- SAN virtualization can split the difference, providing a single method of replication for multiple types of storage arrays and a limited number of management points.