STORAGE Storage Decisions

Storage Challenges Created by a Virtualized Server Infrastructure

Steve Norall
Senior Analyst
Taneja Group
steve@tanejagroup.com

STORAGE Storage Decisions

Agenda

- State of server virtualization
- Four storage challenges created by server virtualization
- How to plan your infrastructure to address these challenges
- Summary

STORAGE Storage Decisions

Virtual Data Center Vision: Compute, Networking and Storage Virtualization

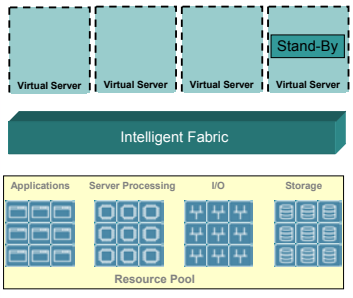

- "Pools" of commonly grouped physical resources
- Dynamic allocations based on application level grouping and usage policies
- Interconnected and controlled through an intelligent interconnect fabric

STORAGE Storage Decisions

What is Server Virtualization?

Without server virtualization:

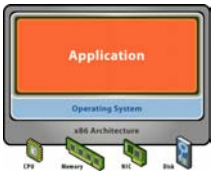

- Single OS image per machine
- Software and hardware tightly coupled
- Running multiple applications on same machine often creates conflict
- Underutilized, inflexible, costly infrastructure

With server virtualization:

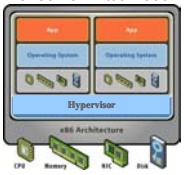

- **Multiple OS images per machine**
- Manage OS and application as single unit by **encapsulating** them into VMs
- VM's are **hardware-independent**: they can be **provisioned & moved** anywhere

STORAGE Storage Decisions

State of Server Virtualization Adoption

- **Server virtualization is moving from test/dev into production**
  - Server consolidation and disaster recovery are two primary usage models
- **85% of all production virtualized environments are attached to a SAN**
  - Historically 100% FC SANs attached, but that is changing…
  - iSCSI and NAS have recently been certified
- **Server virtualization changes how servers talk to storage**
  - Advanced file systems & volume management technologies are embedded into server virtualization
  - CLVMs are required for VM mobility (e.g. VMotion & Live Migration)
  - Single HBA shared across all virtualized guest OSes

STORAGE Storage Decisions

Planning Your Virtualized, Consolidated Infrastructure

- Large scale consolidation of 4-6 VMs per server tightly dependent on storage infrastructure
- Planning of virtualized server infrastructure interlinked to planning storage infrastructure
- Can't optimize one part of the equation without considering the other side
- Storage virtualization is a companion technology to server virtualization

STORAGE Storage Decisions

Four Storage Challenges Exacerbated by Server Virtualization

1. Storage utilization decreased by server virtualization
2. Application performance dependent on storage performance
3. End-to-end visibility from virtual machine through physical device
4. Diagnosibility, tuning, & change management are more difficult

STORAGE Storage Decisions

Challenge #1: Poor Utilization

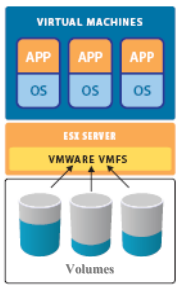

- ESX Server and ESX server management favors large upfront allocations of useable capacity.
- ESX requires additional unwritten space for resume/suspend operations
- Even though all guest operating systems may not yet be configured. And configured guest operating systems may only have written to a small portion of pre-allocated useable capacity.
- Cost of pre-allocated disk capacity can be a major inhibitor to broader deployment of VMware

STORAGE Storage Decisions

Solution: Thin Provisioning

- Thin provisioning only commits physical capacity on written data
- Result: Physical capacity is used more efficiently
- Thin provisioning is found in some virtualization appliances and arrays
- Examples: 3PAR, Compellent, DataCore, NetApp, etc.

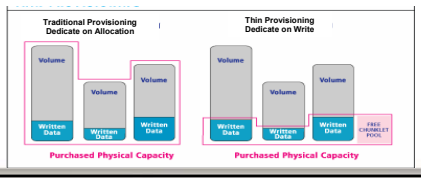

STORAGE Storage Decisions

Challenge #2: Application Performance Dependent on Storage Performance

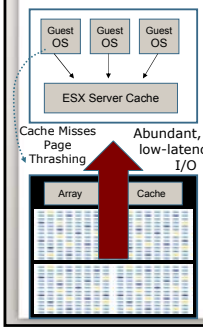

- Consolidate, virtualized server infrastructure creates a unique I/O profile
- Memory is scarcest resource in modern day servers
  - 60% utilized on typical physical server
  - Can't add enough RAM for high scale consolidation
- Greater server RAM cache misses generate much higher read I/O to the storage array including additional paging and page thrashing
- High level of I/O randomness generates significant storage array cache misses
- More trips to disk slow the system and application performance
- Net net: Application performance becomes highly dependent on storage responsiveness

STORAGE Storage Decisions

Solution: Wide Striping & Fast, Low Latency Storage

- Need to optimize storage layout for quick responsiveness to read I/O and page retrieval requests
- Striping data across all available spindles is critical to achieve
  - High IOPS
  - Low latencies
- Advanced storage virtualization crucial for storage performance optimization & layout

STORAGE Storage Decisions

Challenge #3: No End-to-End Visibility from VM to Physical Storage

- WWNs assigned to a single physical server not individual VMs
- Traditional storage tools do not report & track to VM
- Difficult to track back allocations & utilizations from VMs to volumes to physical disk

STORAGE Storage Decisions

Solution: Virtual HBAs & NPIV

- Virtual HBAs solve visibility problem
  - A pool of WWNs are made available for each virtual machine to establish its own connection to the SAN
  - Eliminates need for application to share WWNs
- Uses N-Port ID Virtualization (NPIV)
- Enables SAN best practices, standard management tools
  - Eg. Fabric zoning, LUN masking
- Enables application-based policy management
  - Eg.: Billing, Backup policies
- Improves portability of VMs

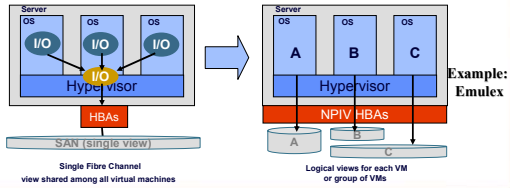

Example: Emulex

STORAGE Storage Decisions

Challenge #4: Diagnosibility, Tuning, Change Management

- Server & storage virtualization drive a more complex SAN fabric
  - Complexity = More management overhead & errors
- VM mobility (e.g. VMotion) creates a changing capacity and performance picture
  - Best laid plans change over time and resources are not used as intended need for tighter change management procedures
  - No historical performance data tied to VMs; only servers
- Traditional management tools perform local capacity planning and performance management
- Optimize capacity and performance for applications, not just servers or storage

STORAGE Storage Decisions

Solution: Cross Domain Mgmt Tools

- New category of tools that span infrastructure domains
  - Network
  - Server
  - Storage
- Focus on optimizing entire infrastructure, not a single domain
- Ideally suited for diagnosing, enforcing change management and optimizing performance holistically
  - Simplified troubleshooting and problem resolution in a virtualized environment
  - Historical performance of VMs
  - Enforce change management for VMotion
- Examples: Akorri, Onaro

Storage Solutions for a Consolidated, Virtualized Infrastructure

1. Advanced storage virtualization in arrays or network
   - Thin provisioning
   - Striping volumes across devices
2. Virtual HBAs (NPIV)
3. Cross-domain management tools

Summary

- Server virtualization causes new storage management challenges and exacerbates pre-existing ones
- Large scale server consolidation requires careful storage infrastructure planning
- Key lesson: Must think holistically across both server and storage domain
  - Don't optimize for one domain at detriment of the other
- Storage virtualization is a key fulcrum in optimizing performance & cost efficiency in a consolidated environment

Questions? Thank You.

Steve Norall
Taneja Group
steve@tanejagroup.com