Thin Provisioning: Keys for Using it to Your Best Advantage

Jeff Byrne & Christine Taylor
The Taneja Group
June 2009

Agenda

- Intro to the Technology
- Why it's Important
- Major Use Cases
- Overall Benefits
- Representative Vendors & Offerings
- Risks & Issues
- Key Take-aways

Thin Provisioning: What It Is, What It Does

- Problem: Over-provisioning to ensure sufficient capacity
- Thin provisioning virtualizes storage for cost-effective capacity management
  - Writes-on-demand keep physical storage allocation to a minimum
  - Physical capacity automatically allocated as needed from a common storage pool

STORAGE Storage Decisions

Thin Provisioning: Example

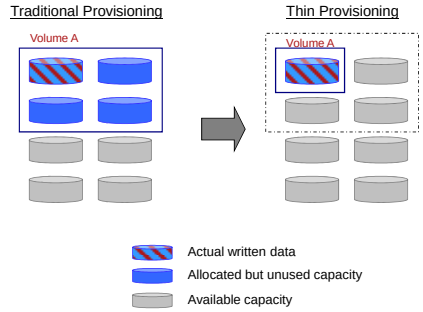

STORAGE Storage Decisions

Not a Panacea

- Requires knowledge about storage infrastructure and present allocation
- Set capacity threshold alerts for common storage pool
- Not all applications will work with or benefit from thin provisioning
  - Avoid apps whose storage needs grow unpredictably, including huge spikes
  - Examples: high-vol databases, scientific apps
  - Some applications require physical allocation

STORAGE Storage Decisions

Major Use Cases

- Provisioning management
- Virtual server storage
- Green data center

STORAGE Storage Decisions

Challenge #1: Over-provisioning

- Users and applications often request more physical storage than they really need
  - Ensure plenty of space for future growth needs
  - Spend allocated budget dollars before they go away
- Storage admins also provision based on conservatively high estimates
  - Quite painful to have to re-size a LUN later
- The result:
  - Poor utilization rates (often in 15-20% range)
  - Inefficient, even wasteful, use of storage capacity
  - Higher costs

STORAGE Storage Decisions

Replication of Unwritten Data

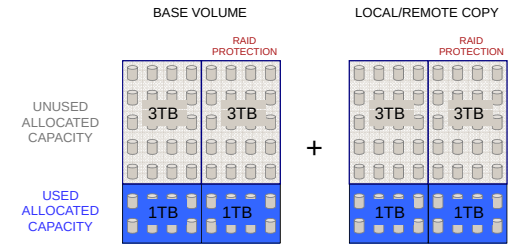

TRADITIONAL PROVISIONING	
16 TB	$480,000

(at $0.03 per MB)

STORAGE Storage Decisions

Solution: Thin Provisioning

- Admins allocate more storage to apps/users than is physically available
  - Example: IT virtually provisions 100 GB to an application while physically provisioning only 10 GB out of the storage pool
- As data needs grow, TP automatically releases additional physical chunks
- When data volumes are RAID protected or replicated, only used and allocated space will consume additional capacity
- Admins can set capacity alerts to enable additional provisioning if needed

STORAGE Storage Decisions

Challenge #2: Poor Utilization in a Virtual Server Environment

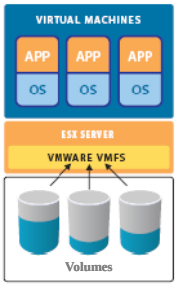

- ESX Server and ESX Server management favors large upfront allocations of usable capacity.
- ESX requires additional unwritten space for resume/suspend operations
- Even though all guest operating systems may not yet be configured. And configured guest operating systems may only have written to a small portion of pre-allocated usable 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 many virtualization appliances and storage arrays commonly used in virtual server environments
- VMware itself is now offering thin provisioning

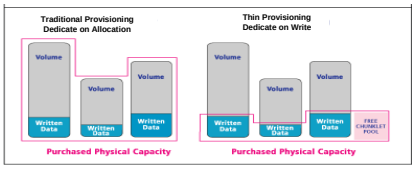

STORAGE Storage Decisions

Challenge #3: High Energy Costs

- Power and cooling in the data center is a huge cost
  - Data centers account for at least ¼ of power and cooling in the entire Facilities budget
  - Average cost of cooling alone is $60 per square foot of the data center
- Storage is a major factor
  - The more spindles you use, the more power and cooling you need

STORAGE Storage Decisions

Solution: Thin Provisioning for the Green Data Center

- Green data center is more than thin provisioning
  - Thin provisioning, data deletion, storage tiering, compression, de-duplication, and hardware technologies
- Work in tandem to reduce power and cooling costs in the data center
  - Thin provisioning is an important component
- Reducing the carbon footprint

STORAGE Storage Decisions

Overall Benefits

- Minimize and economize use of physical storage
- Dramatically increase storage utilization
- Save on CAPEX and OPEX costs
- More easily accommodate storage needs of a mix of applications and users
- Increased application performance (via wide striping)
- Reduced administrative time/costs

STORAGE Storage Decisions

Thin Provisioning Vendors

- Lots of vendors advertise thin provisioning
  - Make sure their product is not just storage virtualization
  - Thin provisioning is a specialized subset of storage virtualization, not a replacement
- A few thin provisioning vendors:
  - EMC, 3PAR, NetApp, FalconStor, Compellent, DataCore, IBM SVC, HDS
  - We'll look briefly at some of these

STORAGE Storage Decisions

3PAR Thin Provisioning

- Early thin provisioning champion; delivers utility computing architecture and highly virtualized storage
- Dedicate-on-write as opposed to dedicate-on-allocation
- Administrators initially provision virtual storage to application
- Physical storage automatically allocated in small amounts for growing application needs
  - On-demand increments avoids the opposite problem of buying large amounts of physical storage to feed the storage pool

STORAGE Storage Decisions

EMC Virtual Provisioning in Control Center

- Symmetrix DMX, CX Clariion, Celerra; works on all tiers including Tier 0 flash storage drives
- Thin device holds application's virtual allocation
- Data device underlies thin devices
  - Virtualized storage automatically striped across multiple spindles to increase performance
  - Acts like separate LUNs for separate management and RAID protection
  - Preserves logical relationships to physical devices; displays in EMC SAN Manager

STORAGE Storage Decisions

NetApp Data ONTAP 7G with FlexVol

- Data ONTAP 7G is NetApp's OS for dynamic workload sharing (FlexShare) and thin systems (FlexVol)
- FlexVol creates flexibly sized LUNs and volumes across a large pool of disks and one or more RAID groups
- Manage from NetApp's policy-based Provisioning Manager
- Practices space reclamation on NAS volumes (CIFS)

STORAGE Storage Decisions

DataCore SANmelody

- Downloadable software; array vendor-neutral
- Space reclamation: Automatically reclaim allocated or thinly provisioned space in background
- Over-subscription: Monitoring services for thin provisioning built into SAN configuration manager interface; issues early alert that common storage pool will run out of disk space
- Hot Swap migration: Ability to non-disruptively migrate thin provisioned data between physical disks in the storage pool

STORAGE Storage Decisions

IBM SAN Volume Controller

- IBM SVC creates block-based virtual storage pools for applications using IBM or non-IBM storage
- Accomplishes thin provisioning with “Space-Efficient Virtual Disks” (SEV)
- Thin provisions any SVC-controlled disk even if disk vendor does not offer thin provisioning

STORAGE Storage Decisions

HDS Universal Storage Platform V

- Positions USP-V as matching storage class to applications
- USP-V can virtualize storage from EMC, IBM, HP and others
- Similar to IBM SVC, this includes thin-provisioning HDS and non-HDS disk attached to its controller

STORAGE Storage Decisions

Risks & Issues

- Potential for running out of space
- Some applications, as well as file systems and databases, are not a good fit
- Mileage may vary, based on vendor approach
- To mitigate these risks and issues:
  - Qualify candidate applications in advance
  - Check out vendor approaches to ensure a fit
  - Decide what type of operational monitoring is required, and make sure it happens

STORAGE Storage Decisions

Key Takeaways

- Thin provisioning:
  - Allocates virtual storage space to suit application requirements
  - Allocates physical storage on an actual use basis
  - Physical allocation should happen automatically from a common storage pool
  - Doesn't work with everything but can bring real benefits to the data center table
  - Good for controlling over-provisioning and energy costs in the data center

STORAGE Storage Decisions

Questions?

jeff.byrne@tanejagroup.com
christine@tanejagroup.com