



Storage Management for Virtual Environments

Jacques Butcher

Senior IT Specialist – Tivoli Storage

jbutcher@au1.ibm.com

Thursday, 31 May 2012

Pulse2012

Meet the Experts. Optimise your infrastructure.

May 31 – June 1


Sheraton on the Park Hotel, Sydney

Agenda

- Introduction to IBM Virtualization Solutions
- Data Protection and Recovery for Virtual Environments
- Management of Virtual Servers and Storage
- Storage Virtualization for Virtual Servers
- Summary
- Evaluations

IT Transformation Roadmap for Virtualized Environments

Physical Consolidation



- Improve utilization
- Reduce costs
- Lower power usage

Improve capacity utilization by as much as 60%, while reducing the power and cooling costs

Advanced Virtual Resource Pools



- Decouple complexity from scale
- Share resources optimally
- Automate workload management
- Incorporate HA & DR

Hands-free operation, eliminate mundane tasks and manual processes and deploy workloads in minutes

Fully virtualized IT with Integrated Service Management



- Sense and respond to workload requirements
- Dynamically move workloads to best-fit infrastructures
- Integrated virtualization management with IT processes

Save time and reduce skill level required for workload provisioning through pre-packaged automation templates

Cloud



- Low cost through economies of scale
- Always on
- Globally available
- Elastic scaling
- Pay for use
- Self-service with rapid provisioning
- Service catalog

Give users the flexibility to request and pay for services they want without the complexities of establishing an IT infrastructure

Comprehensive IBM Virtualization Offerings

Server Virtualization

- System p, System i, System z LPARs, VMware ESX, IBM Smart Business Desktop Cloud
- Virtually consolidate workloads on servers




File and File System Virtualization

- Scale out NAS (SoNAS), Storwize V7000 Unified, DFSMS, IBM General Parallel File System, N-series
- Virtually consolidate files in one namespace across servers



Disk and Tape Storage Virtualization

- SAN Volume Controller (SVC), ProtecTIER
- Industry leading Storage Virtualization solutions



Server and Storage Infrastructure Management

- Data protection with Tivoli Storage Manager and TSM FastBack
- Advanced management of virtual environments with TPC, IBM Director VMcontrol, TADDM, ITM, TPM
- Consolidated management of virtual and physical storage resources



IBM Storage Cloud Solutions

- Smart Business Storage Cloud (SoNAS ACE), IBM SmartCloud Managed Backup
- Virtualization and automation of storage capacity, data protection, and other storage services





Data Protection and Recovery for Virtual Environments

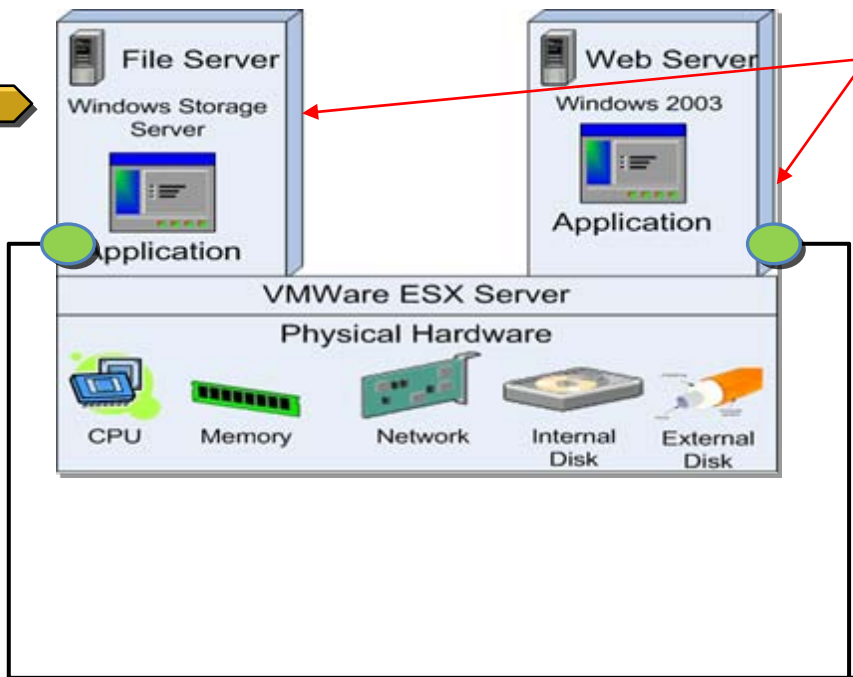
Pulse2012

Meet the Experts. Optimise your infrastructure.


May 31 – June 1

Sheraton on the Park Hotel, Sydney

Traditional (Guest OS) Backup using TSM or TSM FastBack



1. Install TSM or TSM FastBack client inside the guest OS
2. Perform very efficient block level incremental snapshots to FastBack server
3. For longer retention you can migrate the data to TSM



 TSM client, TSM FastBack, or TSM for DB/Mail/SAP/SharePoint agent

Traditional (Guest OS) Backup – Pros and Cons

A large percentage of users are still using this approach. TSM supports many hypervisors (KVM, Hyper-V, VMware, LPARs, Solaris Containers, HP nPartitions, etc) and guest OSs (Windows, Linux, zLinux, z/OS, Solaris, etc). TSM Fastback's Block level incremental forever backup makes it a very good fit for this approach since the Fastback client has very low overhead.

Pros	Cons
Application awareness during backup	Multiple agents
Better recovery for application	Management challenges
CDP protection with TSM FastBack	Lacking VMware integration
Business as usual, use existing management methodologies	



TSM for Virtual Environments (TSM for VE)

Support multiple recovery options from image backup and vStorage API change block tracking (CBT)
New TSM for * (additional component) enhancing the b/a client (Windows only) support to include

- CBT allowing incremental backups (with periodic block-level full backups)
- File/Volume/Disk/Full VM restores from an image backup (multiple OSs are supported)

Added Value


- Single Source Backup
- Change Block Tracking (CBT)
- File level recovery from any OS
- Near-Instant Volume Restore

DP for VMware

Running on Windows proxy*



1 Backup VM



Tivoli Storage Manager
Storage Pools

vStorage API

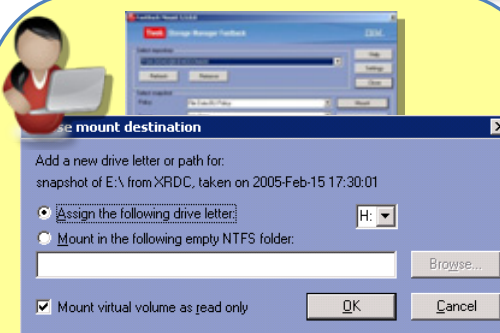
Restore single file
directly to guest (or any
other target)

3

Mount image
directly from TSM
disk storage pool,
expose it locally or
using an iSCSI
target interface

2

* Proxy server can be a physical or
virtual machine



TSM for Virtual Environments (TSM for VE) – Pros and Cons

Supported **only on** VMware, this approach tries to combine the benefits of traditional and host based approaches by providing an API to talk to the console and move the data through the proxy server. This is the recommended approach by VMware.

<http://www-01.ibm.com/support/docview.wss?uid=swg21394300>

Pros	Cons
Saves money <ul style="list-style-type: none">•no physical proxy with disk required•saves valuable administration time	Questionable application integration (VMware does trigger VSS for windows guests)
Reduces downtime	Requires regular full backups (VMware)
“LAN Free” backup (needs physical proxy)	Additional disk storage Requirements on TSM server
Offloads backups to proxy server	
Flexibility – supports both file level and image level backup.	
Utilizing VMware API including Changed Blocks API	
Recommended by VMware	



TSM for Virtual Environments (TSM for VE)

- VMware creates a temporary snapshot of a virtual machine:
 - VMware device driver flushes buffers on filesystem within VM
 - MS VSS Provider interfaces with OS, File System and VSS Writers (Applications and system components)
 - Snapshot of VM is taken (.VMDKs are frozen)
 - Redo log is created – all writes are redirected to redo log
 - Virtual machine is backed up
 - Virtual Machine Snapshot is deleted, causing Redo log data to be applied to the original .VMDKs
 - Backup complete
- Every step involves I/O
- The faster the backup, the sooner performance impacting snapshots can be removed!

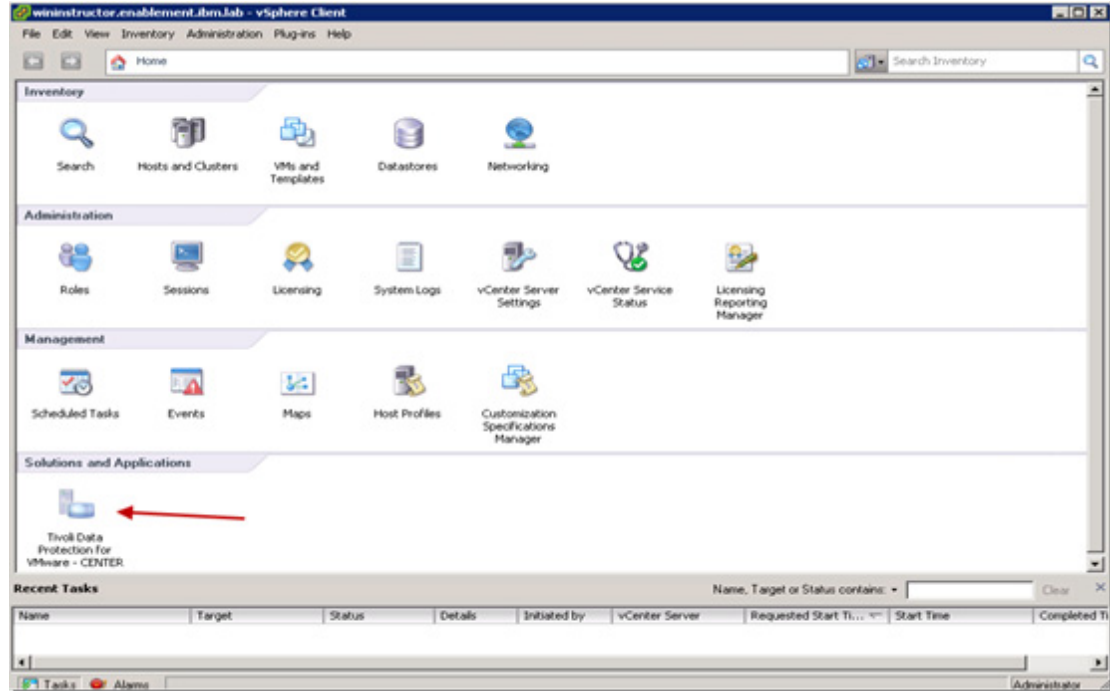
TSM for VE Data Movers

- **Multiple Datamovers on multiple machines**
 - Run in parallel on different machines
 - Can be Virtual Machines
 - Can be Physical Machines
 - Support multiple vSphere Datacenters
- **Multiple Datamovers on one machine**
 - Run in parallel on the same machine
 - Concurrent backup
 - Can have different schedules
- **Plug-in Integration**
 - Can be configured from the same Plug-in
 - Can be scheduled and managed from the same Plug-in



TSM for VE vCenter Plug-In

- Integrated into vCenter GUI
- Ad-hock and schedule backup & restores
- Reporting
- Configuration



TSM for VE - vCenter Plug-In

The screenshot shows the vSphere Client interface with the following components:

- Hosts Panel:** A tree view showing the hierarchy: vCenter > FCM_VE_Class > vsphere03.enablement.ibm. > Student_10_Hull through Student_9_Hull.
- Navigation Tabs:** Getting Started, Summary, Backup, Restore, Reports, Configuration.
- Main Content Area:**
 - Section Header:** "What is IBM Tivoli Storage Manager for Virtual Environments?"
 - Description:** "IBM Tivoli Storage Manager for Virtual Environments is a data protection solution designed specifically for VMware. The Tivoli Storage Manager data mover performs block-level backups and restores of virtual machines to a Tivoli Storage Manager server. These single source backups can also be used for file restore or instant restore of a volume."
 - Diagram:** A network diagram showing a Vcenter Server Client connected to two Hosts in a Datacenter. The Hosts are connected to a Tivoli Storage Manager Server and a Tivoli Storage Manager Agent.
 - Section Header:** "What do you want to do?"
 - Actions:**
 - Define a backup task ...
 - Initiate a restore ...
 - View active task status...
 - Explore Further:** A box containing links for "Understanding backups" and "Understanding restores".
- Footer:** "Recent Tasks" and a search bar "Name, Target or Status contains: -" with a "Clear" button.

TSM for VE - vCenter Plug-In Scheduler

Schedule a Backup

- Welcome
- General
- Source
- Destination
- Schedule
- Repetition**
- Summary

Repetition

* Date and time of the first backup:
10/25/2011 3:00 PM

Back up weekly

Back up every
 months

Back up on the following days of the week

Monday Tuesday Wednesday Thursday

Friday Saturday Sunday

Newly added virtual machines are included in this backup task

< Back Next > Cancel



TSM for VE - Data Movers

The screenshot shows the vSphere Client interface for a TSM agent named 'repmonserv'. A 'Schedule a Backup' wizard is open, with the 'Destination' step selected. The wizard includes a sidebar with steps: Welcome, General, Source, Destination, Schedule, and Summary. The 'Destination' step contains the following text: 'The destination you choose determines where the VMware snapshot is stored on the Tivoli Storage Manager.' and 'Select the Tivoli Storage Manager data mover node that will run the backup. Pick a data mover node that is not used by another process to improve backup performance.' Below this text is a dropdown menu with three options: 'VC1_DC1_DM1', 'VC1_DC1_DM1', and 'VC1_DC1_DM2'. The 'VC1_DC1_DM1' option is currently selected. At the bottom of the wizard are buttons for '< Back', 'Next >', and 'Cancel...'. Below the wizard is a 'Recent Tasks' table.

Name	Target	Status	Details	Initiated by	vCenter Server



TSM for VE Business Benefits

- **Reduces costs**
 - No additional hardware required
 - Retention management integrated into Tivoli Storage Manager policies
 - Simplified agent management with the centralized IBM Tivoli® Storage Manager console and one agent supports multiple VMs

- **Reduces risk**
 - Automated discovery of new VMs ensuring you environment remains protected



TSM for VE - Business Benefits (continued)

- **Improves RTOs and RPOs**

- Leverages vStorage APIs for Data Protection (VADP) and Change Block Tracking (CBT)
- Non-disruptive, single-pass, content-aware, block-level backup enabling faster, more frequent protection for virtual machines
- Support for LAN-free data transfer from the VMware server's storage to the backup server - preserving bandwidth for other uses
- Flexible recovery options: file, volume, VM image (BMR)
- Near-instant restore of files and disk volumes (Windows and Linux)

- **Improves production server response times**

- 'Near Zero Impact Backup' - Offload the backup workload from virtual machines and production VMware ESX hosts to vStorage backup proxy servers





FlashCopy Manager (FCM) for Virtual Environments

Pulse2012


Meet the Experts. Optimise your infrastructure.

May 31 – June 1

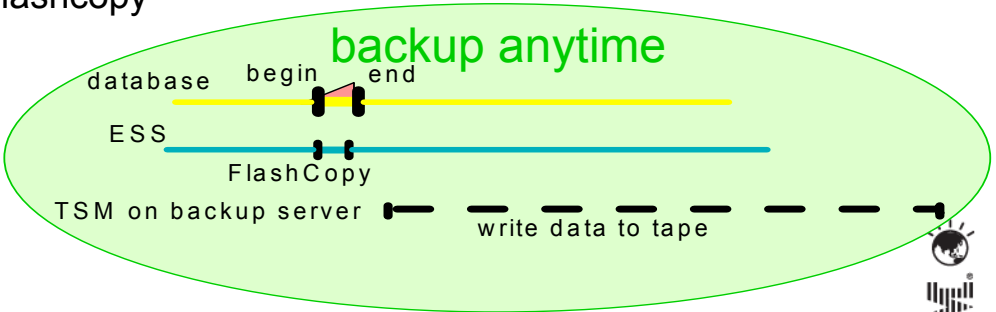
Sheraton on the Park Hotel, Sydney

Hardware Based – Triggering Hardware Based Snapshots

Production



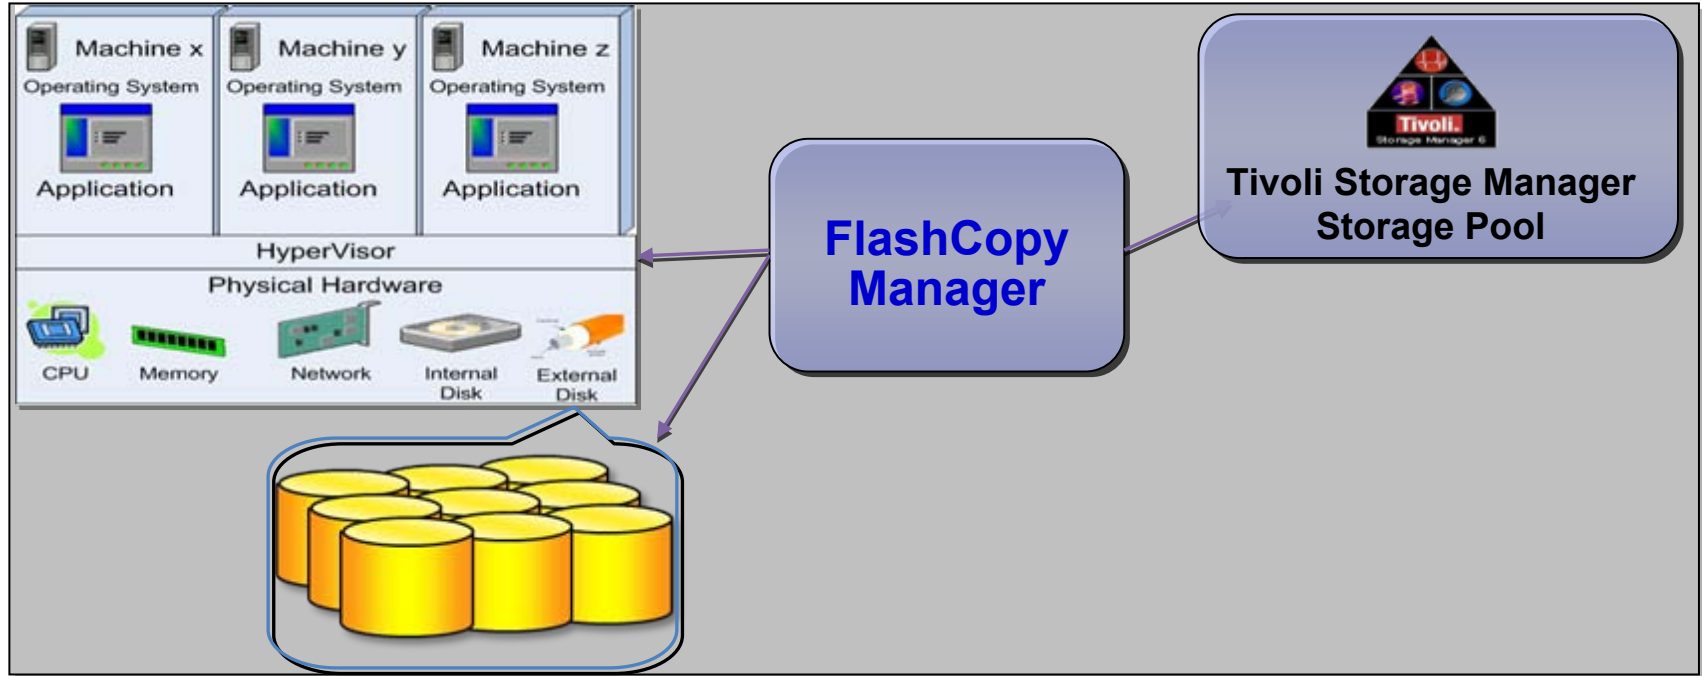
Offload Backup with Flashcopy



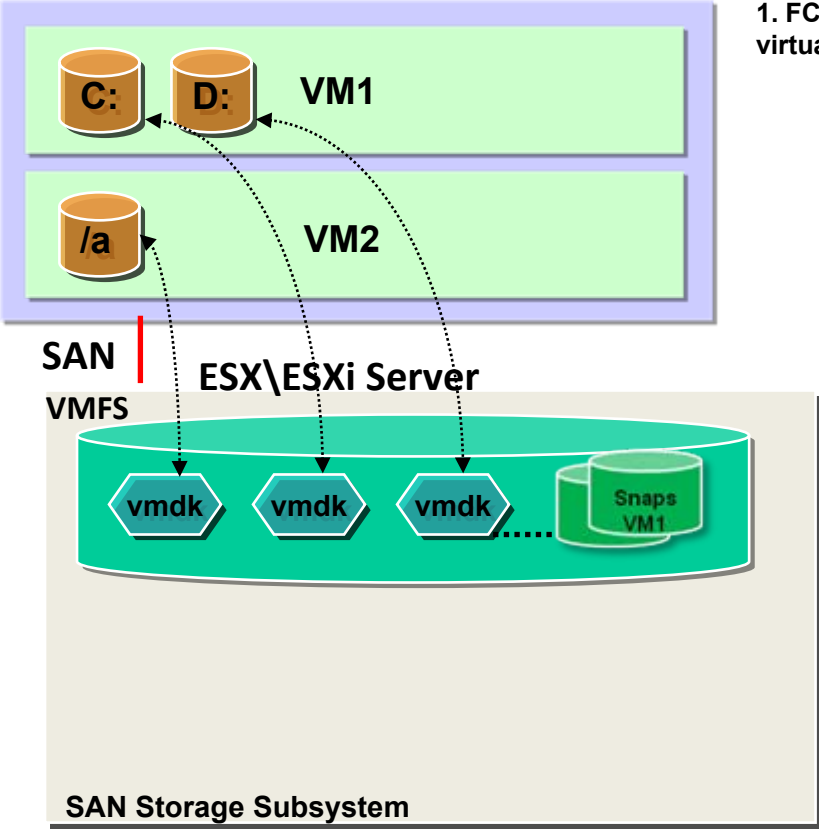
Tivoli Storage FlashCopy Manager (FCM)

Hardware Based – Triggering Hardware Based Snapshots

The “triggering agent” can reside in multiple locations, the host or a proxy server



Full-VM Backup with FCM Leveraging vStorage API




1. FCM initiates a software snapshot of virtual guest volumes (vSphere API)

Linux Proxy Server
(physical or virtual machine)
FlashCopy Manager for VMware




TSM Server

Full-VM Backup with FCM Leveraging vStorage API



1. FCM initiates a software snapshot of virtual guest volumes (vSphere API)


2. FCM determines which LUN(s) are associated with virtual machines

Linux Proxy Server
(physical or virtual machine)

FlashCopy Manager for VMware



TSM Server

Full-VM Backup with FCM Leveraging vStorage API



1. FCM initiates a software snapshot of virtual guest volumes (vSphere API)

2. FCM determines which LUN(s) are associated with virtual machines

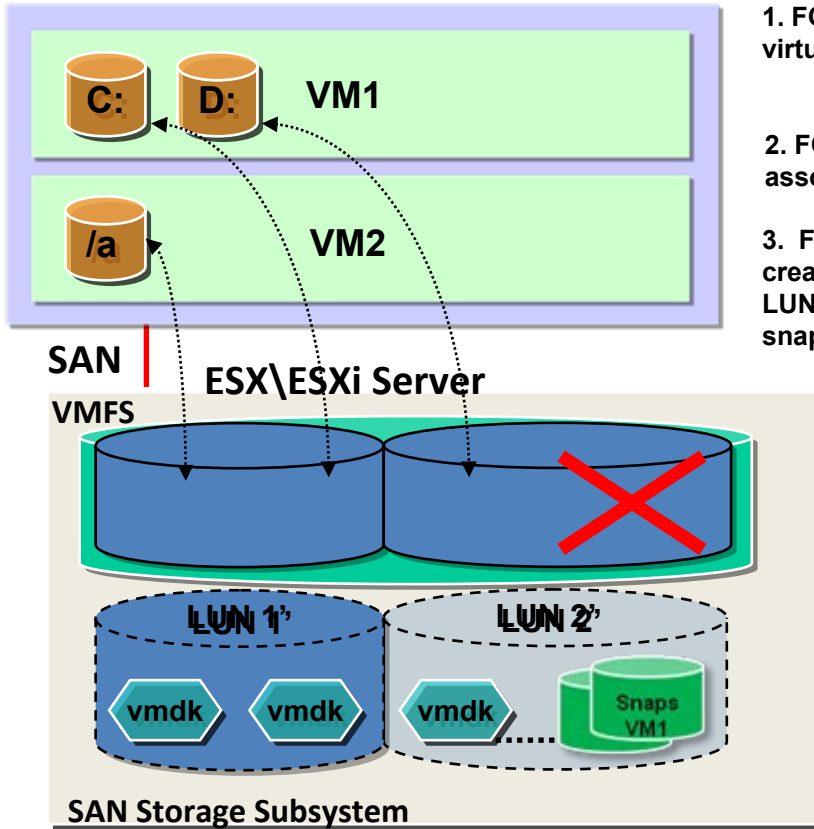
3. FCM invokes hardware FlashCopy to create a persistent snapshot copy of the LUN(s) hosting the .vmdk and software snapshot

Linux Proxy Server
(physical or virtual machine)
FlashCopy Manager for VMware



TSM Server


Full-VM Backup with FCM Leveraging vStorage API



1. FCM initiates a software snapshot of virtual guest volumes (vSphere API)

2. FCM determines which LUN(s) are associated with virtual machines

3. FCM invokes hardware FlashCopy to create a persistent snapshot copy of the LUN(s) hosting the .vmdk and software snapshot

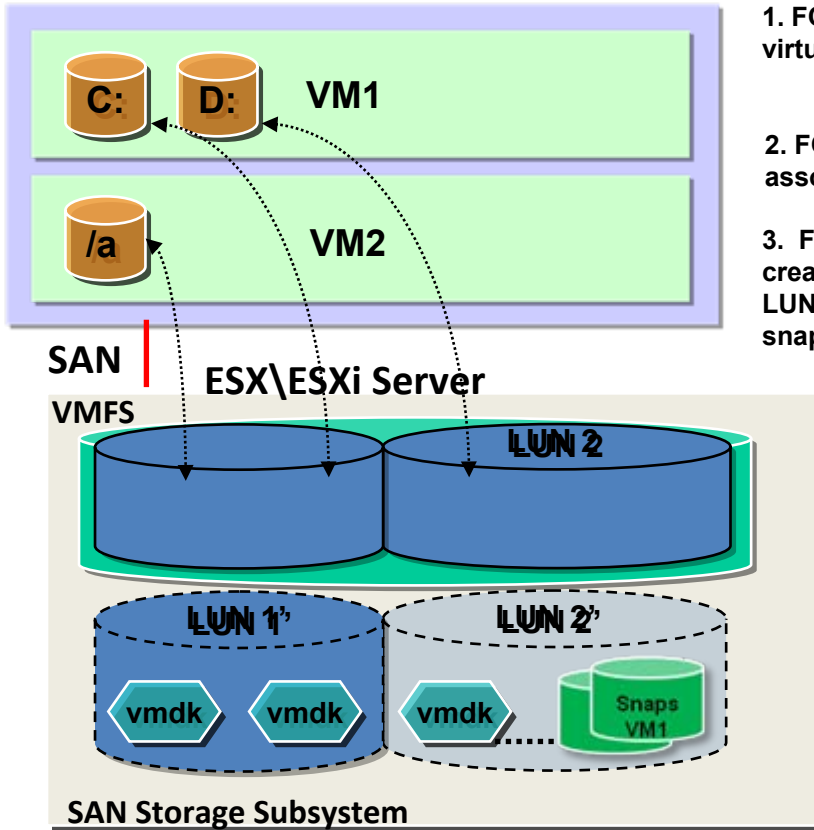
Linux Proxy Server
(physical or virtual machine)

FlashCopy Manager for VMware

4. Hardware snapshot is persisted for use as source for recovery operation, software snapshots are deleted.



TSM Server


Full-VM Backup with FCM Leveraging vStorage API



1. FCM initiates a software snapshot of virtual guest volumes (vSphere API)


2. FCM determines which LUN(s) are associated with virtual machines

3. FCM invokes hardware FlashCopy to create a persistent snapshot copy of the LUN(s) hosting the .vmdk and software snapshot



4. Hardware snapshot is persisted for use as source for recovery operation, software snapshots are deleted.

5. FCM optionally creates additional copy of data on TSM server (Objects saved in common data format to enable individual file recovery using Recovery Agent Mount)



TSM Server

FlashCopy Manager 3.1 for VMware – Features at a glance

- Supports ESX / vSphere 4.1 and 5.0
- Off-host (proxy based) hardware snapshot backups with VMFS datastore granularity
 - Scheduled or ad-hoc execution
 - No need to deploy OS specific agents to each VM
- Optional offloaded backup to TSM with VM granularity
 - Supporting block level incremental backups with VMware Changed Block Tracking
- Restore of individual VMs from
 - a) from a Hardware snapshot of a VMFS datastore
 - b) from offloaded TSM backup to original or alternative VMFS datastore under original name or new one

FlashCopy Manager 3.1 for VMware – Features at a glance

- User Interface
 - vCenter client integration (plug-in web GUI)
 - Command Line Interface (on proxy node)
- Reporting and Monitoring
 - summary view with drill down for details
 - backup & restore stats
 - Managed capacity report
- Policy based management of local snapshot backups

FCM – Application Aware Near-Instance Hardware Based

This approach requires a level of integration and coordination between different components:

1. Underlying disk subsystem
2. Hypervisor (VMware)
3. Applications within the guest OS

Pros

Very quick and efficient HW snapshots

No resource consumption on guest or host

Data can be moved to TSM (offloaded backups)

Cons

HW snapshot will include all the virtual disks that reside on the same LUN

Hard to coordinate the HW Snapshot with the application consistency

Recovering the LUN will recover all of the virtual disks on the same LUN



Management of Virtual Server Storage

Pulse2012

Meet the Experts. Optimise your infrastructure.

May 31 – June 1

Sheraton on the Park Hotel, Sydney



What Needs to be Managed?


- **Servers**

- ESX servers
- VM images
- Applications
- Databases
- File Systems
- Volume Managers
- Host Bus Adaptors
- Virtual HBAs
- Multi-Path Drivers



- **Network Components**

- Switches, Hubs, Routers
- Virtual Devices
- Intelligent Switch Replication



- **Storage Components**

- Volume mapping / virtualization
- Storage Array Provisioning
- VMFS, NAS Filers, SoNAS
- Tape Libraries

How Does it Need to be Managed?

- **Discovery**

- Topology Views
- Asset Management

- **Configuration Management**

- Provisioning
- Optimization
- Problem Determination



- **Performance Management**

- Bottleneck Analysis
- Load Balancing

- **Monitoring, Alerting, and Reporting**

- Asset/Capacity/Utilization
- Accounting/Chargeback
- Performance/Trending
- Problem Reports
- Storage and Data Analysis



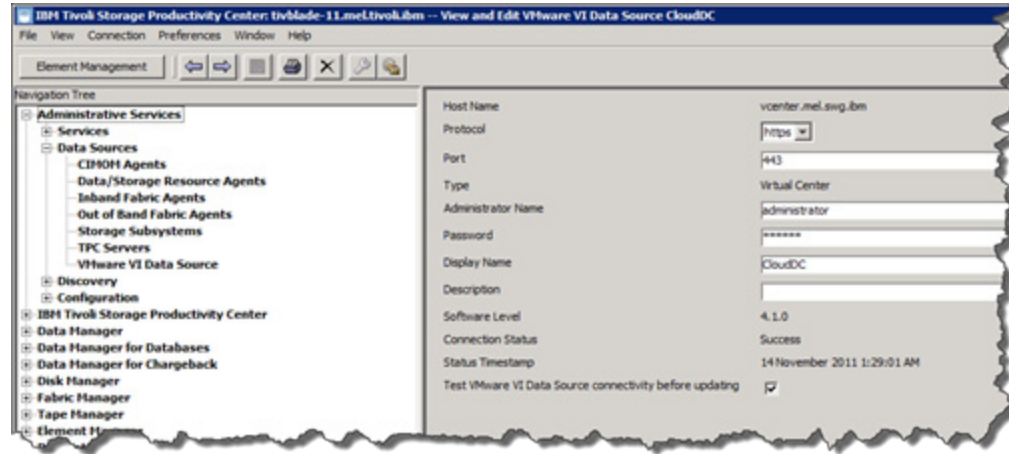
As we are dealing with a network, we need to work with the end-to-end network configuration, not just the individual components

Current TPC Virtual Server Management Capabilities

- TPC provides advanced management for virtual server and storage environments:
 - **Discovery:** ESX server, VM Guest OS images, VMFS, storage and which VM has storage allocated from where
 - **Topology and Visualization:** Hypervisor views including drill down to show all VM images, end to end correlation of SAN storage to ESX server and VM guests
 - **Monitoring and Reporting for ESX server / VM guests:** health status and monitoring, asset reporting, capacity utilization (total, free, used), ..
 - **Problem Determination and root cause analysis of storage problems:** assistance discovering the 'real' problem in a virtual world
 - **Storage Provisioning:** from any storage array to ESX server

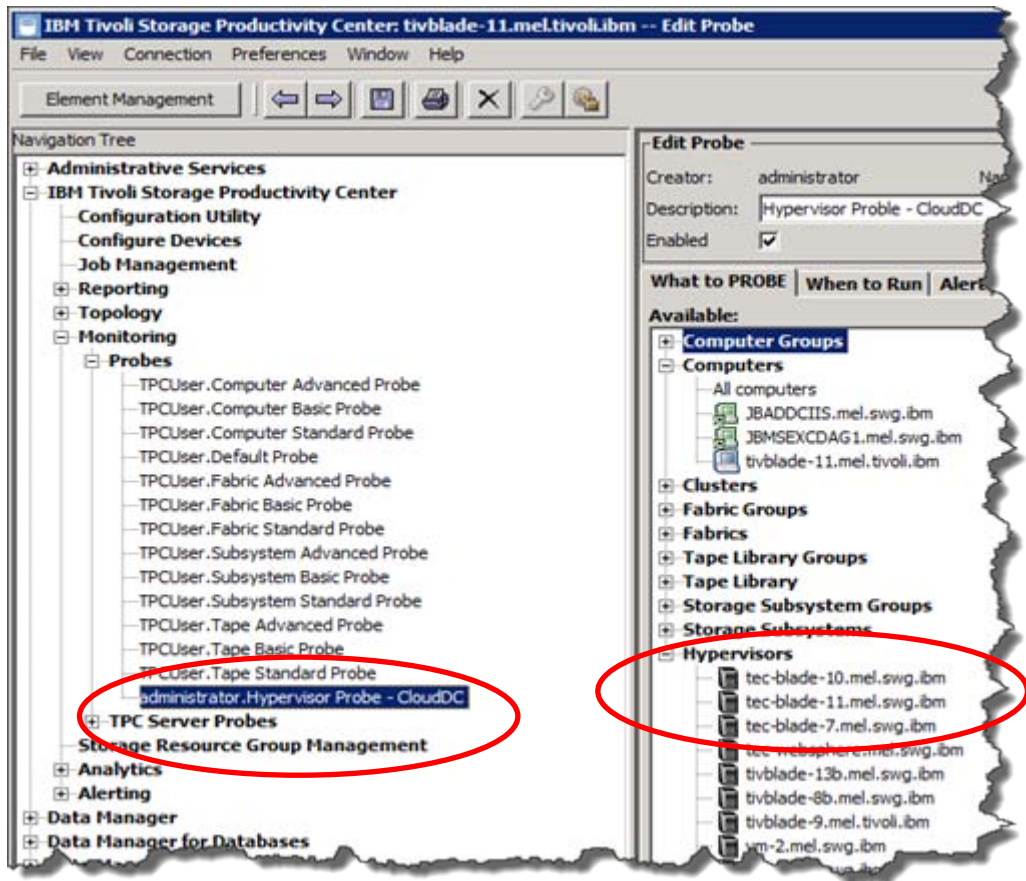
Adding VMware Data Source

- Download vCenter Certificate
- Import Certificate
- Add VC credentials as a data source
- Run Discovery
- Run Probe



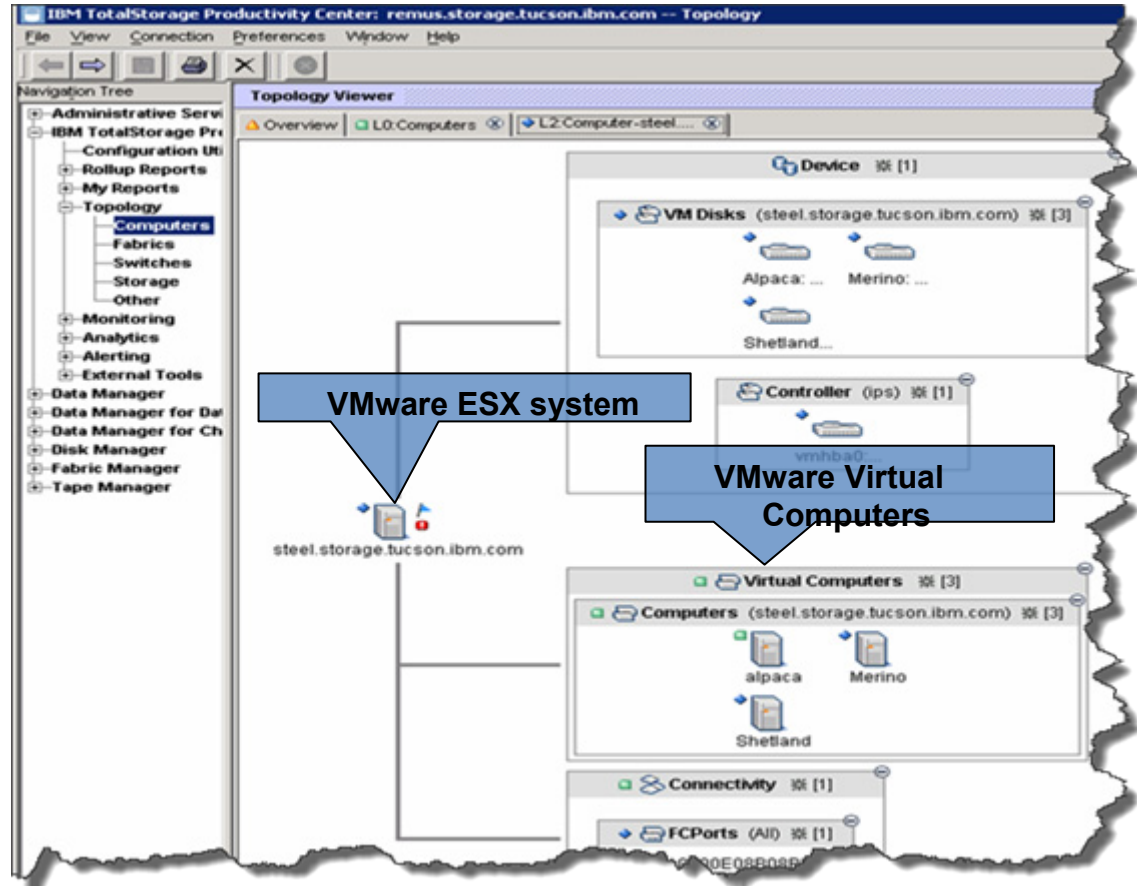
VMware Discovery

- Perform Discovery operation against Hypervisor systems.
- “Hypervisor” added to the list of entities that can be selected for Probes.
- The virtual machines that have the TPC for Data agent installed will be listed as computer systems in the probe definition screen for systems.



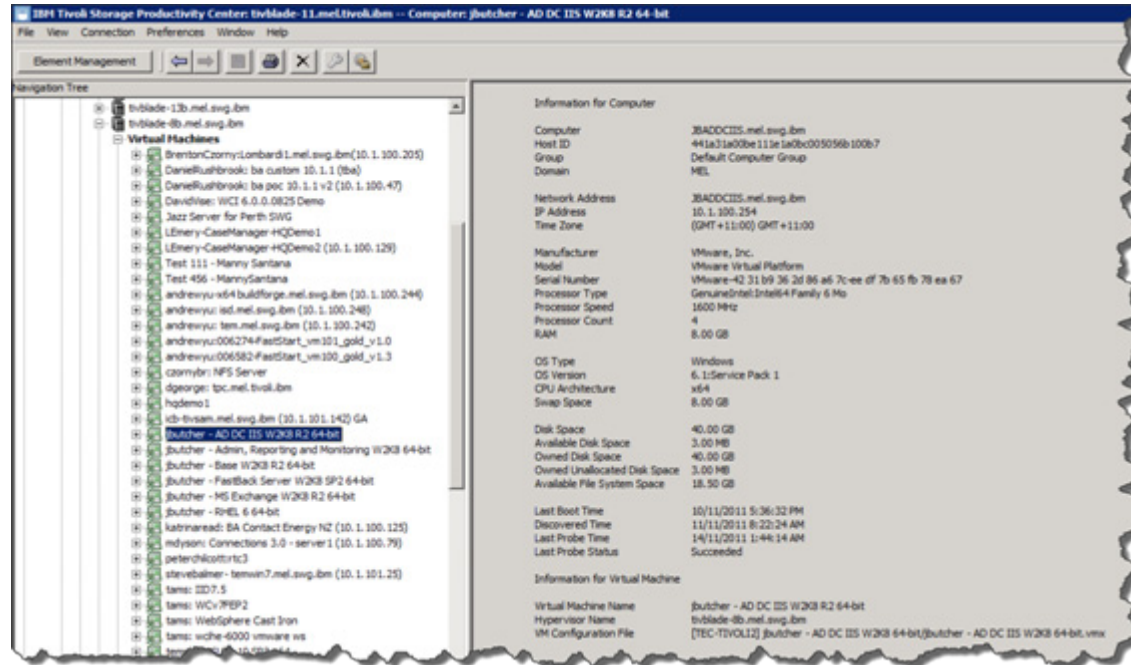
VMware Topology and Visualization

- Discover and report the logical aspects of the VMware environment:
 - *VMware virtual machines and mapping to the host physical machine*
 - *Storage resources used by the ESX server*
- For detailed information on the VMware virtual machine, a TPC Storage Resource Agent (SRA) need to be installed within the guest



VMware Guest OS Properties

- Detailed Asset reports about VMware virtual guests with the TPC SRA installed



VMware Capacity Utilization Report

- Detailed Capacity Reports for VMware ESX System and VMware virtual guests with the TPC SRA installed

The screenshot displays the IBM Tivoli Storage Productivity Center interface. The main window shows a table of disk space utilization for various computers. The table has columns for Computer, Disk Space, Available Disk Space, Owned Disk Space, Owned Available Disk Space, Non-Fibre Channel Attached Disk Space, Non-Fibre Channel Attached Available Disk Space, Fibre Channel Attached Disk Space, and Fibre Channel Attached Available Disk Space. The data is organized by computer, showing both total disk space and available space, as well as the breakdown by channel type (Non-Fibre Channel and Fibre Channel).

Computer	Disk Space	Available Disk Space	Owned Disk Space	Owned Available Disk Space	Non-Fibre Channel Attached Disk Space	Non-Fibre Channel Attached Available Disk Space	Fibre Channel Attached Disk Space	Fibre Channel Attached Available Disk Space
trvblade-11.mel.tivoli.ibm	68.37 GB	2.84 MB	68.37 GB	2.84 MB	68.37 GB	2.84 MB	0	0
tec-blade-7.mel.svg.ibm	7.88 TB	42.90 MB	7.88 TB	42.90 MB	68.37 GB	80.00 KB	7.81 TB	0
trvblade-10.mel.svg.ibm	10.54 TB	55.02 MB	10.54 TB	55.02 MB	67.05 GB	80.00 KB	10.47 TB	0
tec-blade-11.mel.svg.ibm	10.57 TB	30.05 GB	10.57 TB	30.05 GB	67.99 GB	80.00 KB	10.50 TB	0
tec-blade-10.mel.svg.ibm	11.52 TB	60.38 MB	11.52 TB	60.38 MB	67.99 GB	80.00 KB	11.45 TB	0
vm-2.mel.svg.ibm	8.85 TB	48.26 MB	8.85 TB	48.26 MB	33.78 GB	80.00 KB	8.81 TB	0
JPMSEXCDAG1.mel.svg.ibm	140.00 GB	6.00 MB	140.00 GB	6.00 MB	140.00 GB	6.00 MB	0	0
JBADCCIS2.mel.svg.ibm	40.00 GB	3.00 MB	40.00 GB	3.00 MB	40.00 GB	3.00 MB	0	0
JPMSSQL.mel.svg.ibm	140.00 GB	6.00 MB	140.00 GB	6.00 MB	140.00 GB	6.00 MB	0	0



Storage Virtualisation for Virtual Servers

Pulse2012

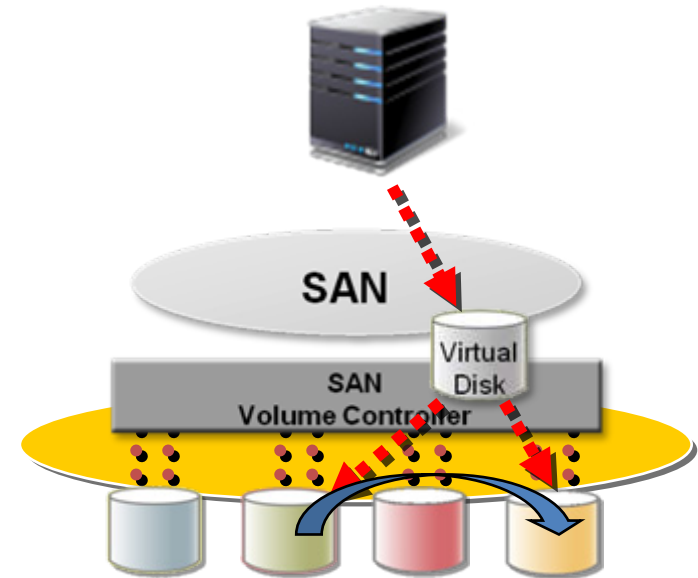
Meet the Experts. Optimise your infrastructure.

May 31 – June 1

Sheraton on the Park Hotel, Sydney

Storage Virtualisation – SAN Volume Controller (SVC)

- Only the SAN Volume Controller seen by the storage disk arrays
 - No advanced function software licensing required on the storage controller
 - Simply provision all the storage to the SVC
 - **Replacing storage does not require changes to the host (no downtime)**
 - Allows thin provisioning, grow your storage only when required




SVC storage virtualization is a perfect match for virtual server environments

Optimized Storage Resource Utilization


Traditional SAN

- Shared physical network
- Limited capacity sharing
- Capacity purchased for, and owned by individual processors
- Poor capacity utilization



SAN Volume Controller


- Hosts own “virtual” disks
- Capacity can be more easily reallocated
- Capacity purchases can be deferred until the physical capacity of the SAN reaches a trigger point.



Improved Application Availability - Migration

Traditional SAN

1. Establish host connections to new storage (reboot may be required)
2. Stop the application
3. Decommission traditional storage
4. **Move data**
5. Restart application



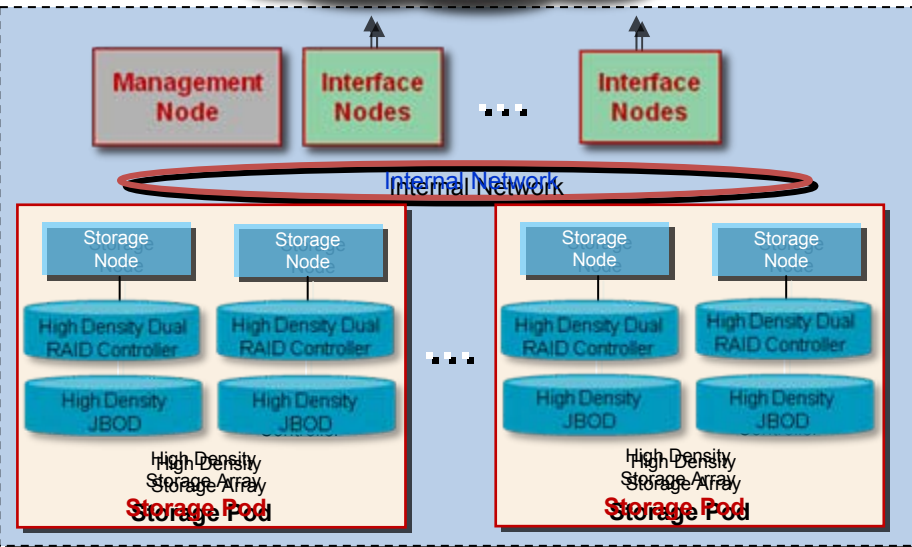
SAN Volume Controller

1. Allocate new storage to SVC
2. Assign to storage pool
3. Move data

No downtime - Host systems and applications are not affected



IBM Scale Out NAS – System Managed Storage in a Box



- Enterprise class solution for IP based filesystem storage (NFS, CIFS, FTP, ..)
- One global repository for application and user files: >1B files per file system, 256 filesystems per SoNAS, simplified management of PBs of storage
- Extreme performance (near linear aggregate throughput) and extreme capacity scaling
- Work load and data is evenly distributed across all nodes and disk pools, eliminating hot spots
- Policy based tiered storage - high-performance SAS and high-capacity SATA HDD's
- Provision, monitor, report, chargeback by application, user, department, etc
- Accelerated backup, HSM and recovery by TSM

Can deploy as private or public cloud



Summary

Pulse2012

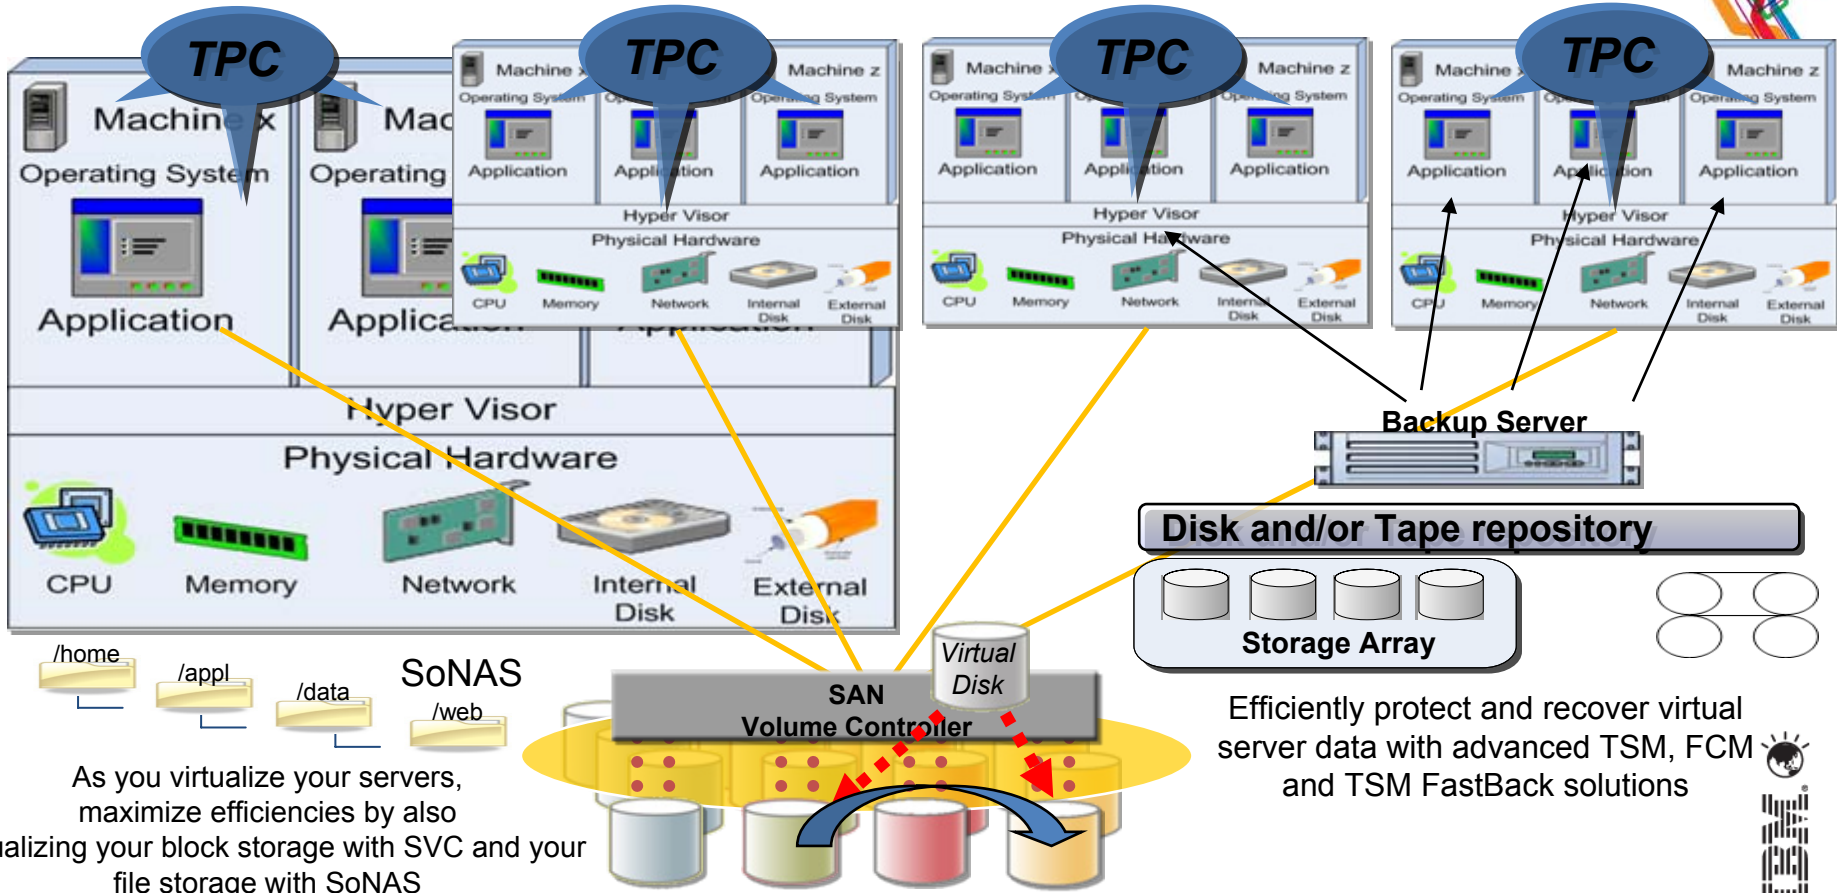
Meet the Experts. Optimise your infrastructure.

May 31 – June 1

Sheraton on the Park Hotel, Sydney



Discover, monitor, report, & Provision Virtual Server & Storage Environments with TPC



/home /appl /data /web
SoNAS

As you virtualize your servers, maximize efficiencies by also virtualizing your block storage with SVC and your file storage with SoNAS

Efficiently protect and recover virtual server data with advanced TSM, FCM and TSM FastBack solutions

Trademarks and disclaimers

© Copyright IBM Australia Limited 2012 ABN 79 000 024 733 © Copyright IBM Corporation 2012 All Rights Reserved. TRADEMARKS: IBM, the IBM logos, ibm.com, Smarter Planet and the planet icon are trademarks of IBM Corp registered in many jurisdictions worldwide. Other company, product and services marks may be trademarks or services marks of others. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

The customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Prices are suggested U.S. list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Photographs shown may be engineering prototypes. Changes may be incorporated in production models.

