

BCT4539

VMware Consolidated Backup Technology: Today and Future

Christian Czezatke

Staff Engineer, VMware



VMWORLD 2006

Overview

- Classic vs. VM Aware Backup
- ESX Snapshots
- COW Files and COW Hierarchies
- Offloaded Backup with VCB
- What VCB is Not
- VCB and 3rd Party Backup Software
- What's Cooking in R&D

Overview

- Classic vs. VM Aware Backup
- ESX Snapshots
- COW Files and COW Hierarchies
- Offloaded Backup with VCB
- What VCB is Not
- VCB and 3rd Party Backup Software
- What's Cooking in R&D

Classic Backup

- Treat Virtual Machines like Physical Servers
 - > Maintain backup agent in every VM
 - > Resource intensive (network, CPU and IO load on ESX server)
 - > Backup going across LAN

VM-Aware Backup

- Backup is based on snapshots of VMs
- Export entire Virtual Machines
 - > Image level backup
 - > Efficient DR (full VM restore)
 - > Often done through COS, either direct or indirect
 - > Resource intensive, load on COS

Overview

- Classic vs. VM Aware Backup
- ESX Snapshots
- COW Files and COW Hierarchies
- Offloaded Backup with VCB
- What VCB is Not
- VCB and 3rd Party Backup Software
- What's Cooking in R&D

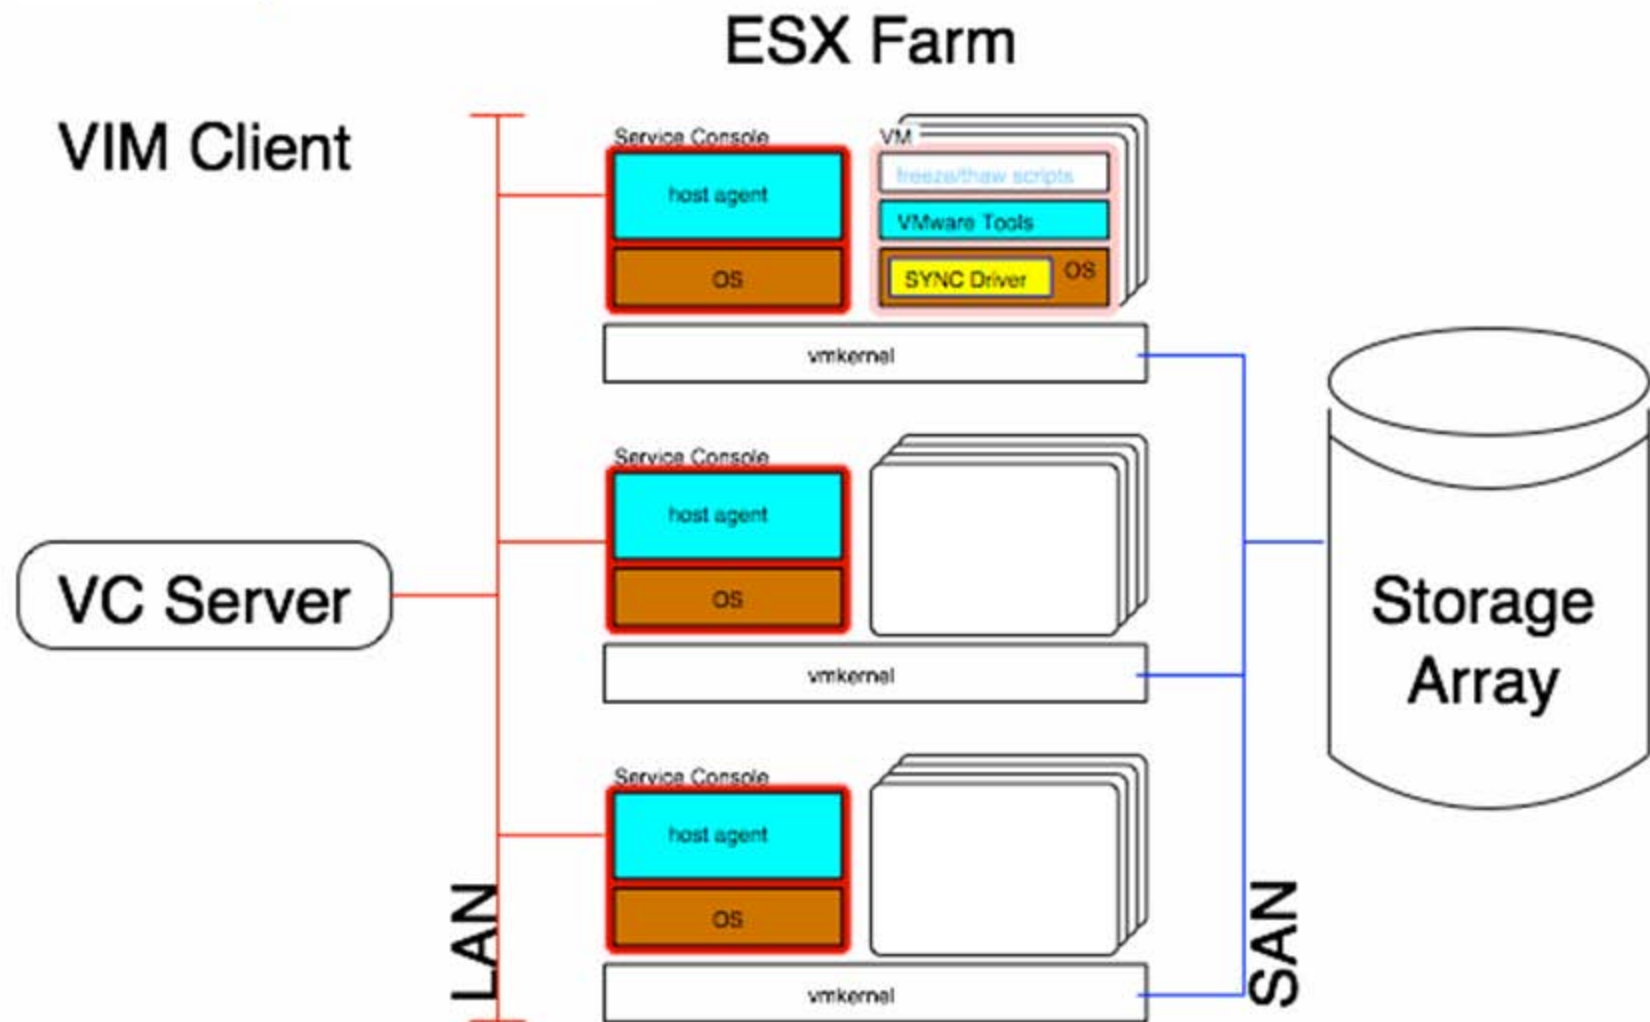
VM Snapshots in VI3 - What's New...

- VI3 introduced enhanced snapshot capabilities for ESX
 - > Snapshot is now VM property, rather than disk property
 - > Snapshots can optionally be quiesced for better data consistency
 - > VIM SDK provides VM snapshot functionality, used by VCB and backup applications

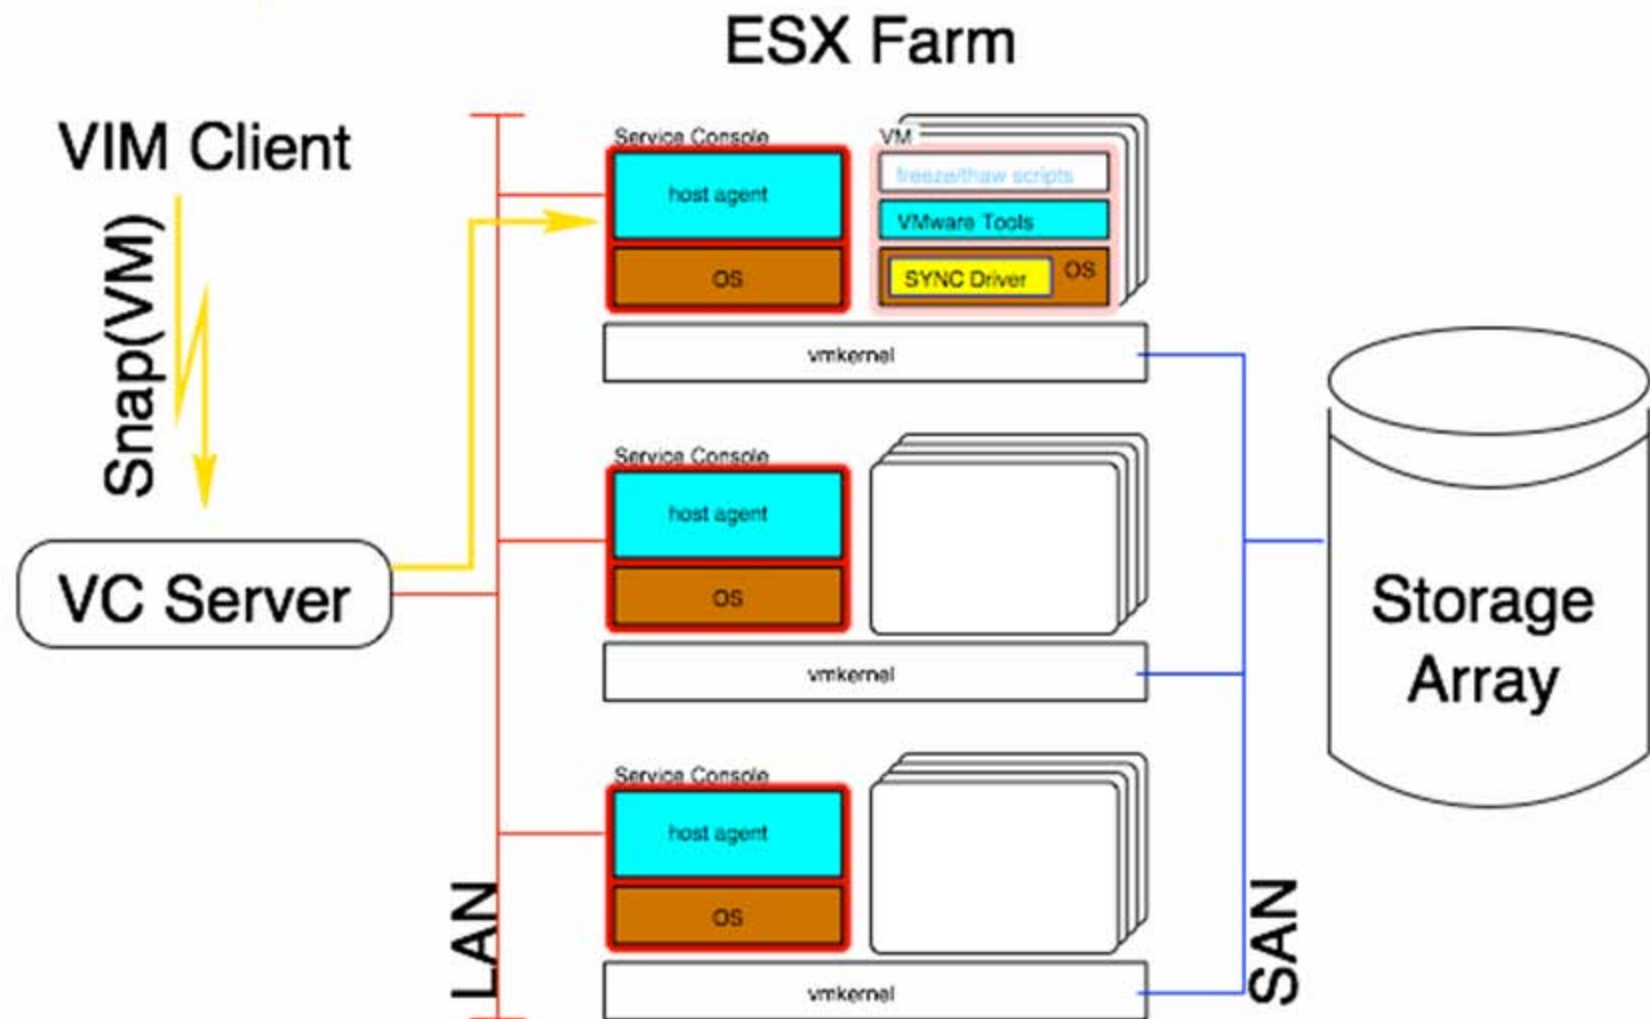
VM Snapshots in VI3

- VI3 VM snapshots involve many components
- VirtualCenter Server
 - > Locates ESX Server currently running VM.
 - > Forwards Request to host agent on ESX
- host agent, drives process on local ESX machine
- VMware Tools, provides quiescing in the guest
 - > Hooks for custom pre-freeze, post-thaw scripts
 - > SYNC driver on Windows XP, 2000, 2003
- ESX vmkernel COW layer ("redo logs")
 - > makes snapshot data immutable

VM Snapshots in VI3

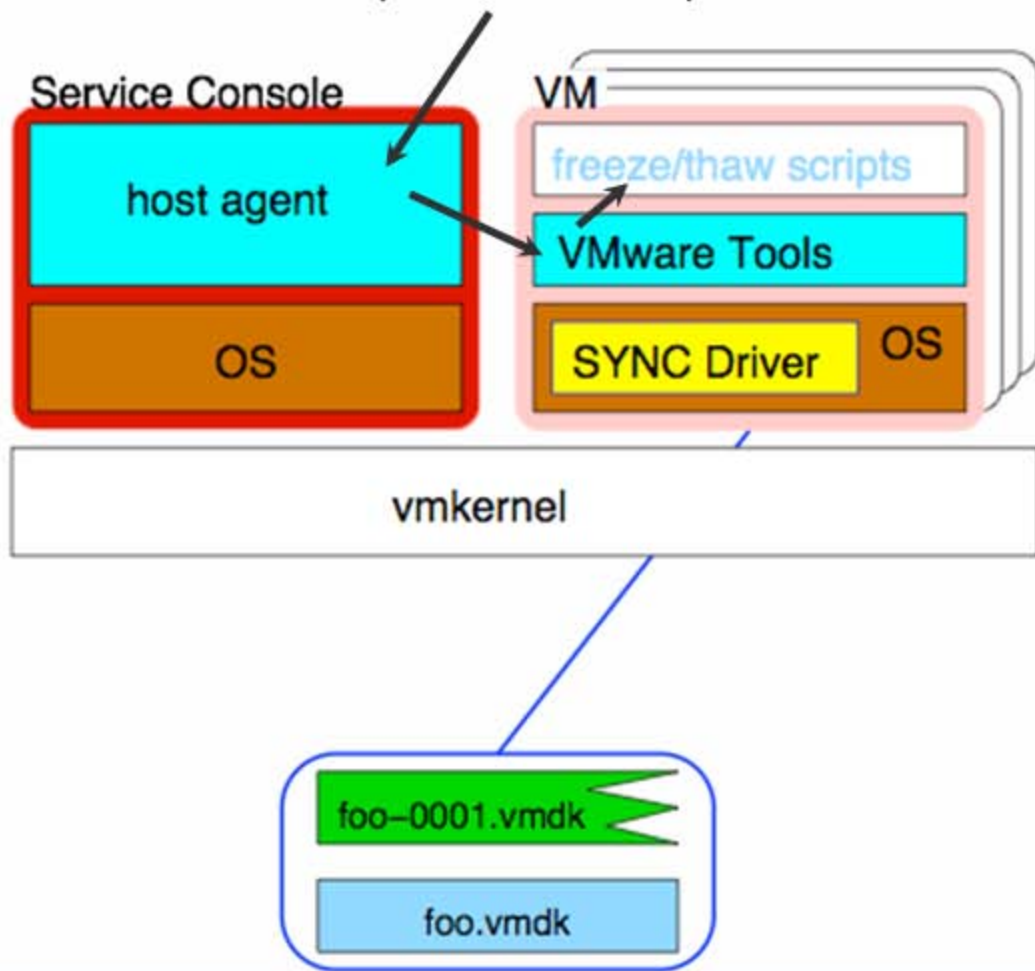


VM Snapshots in VI3



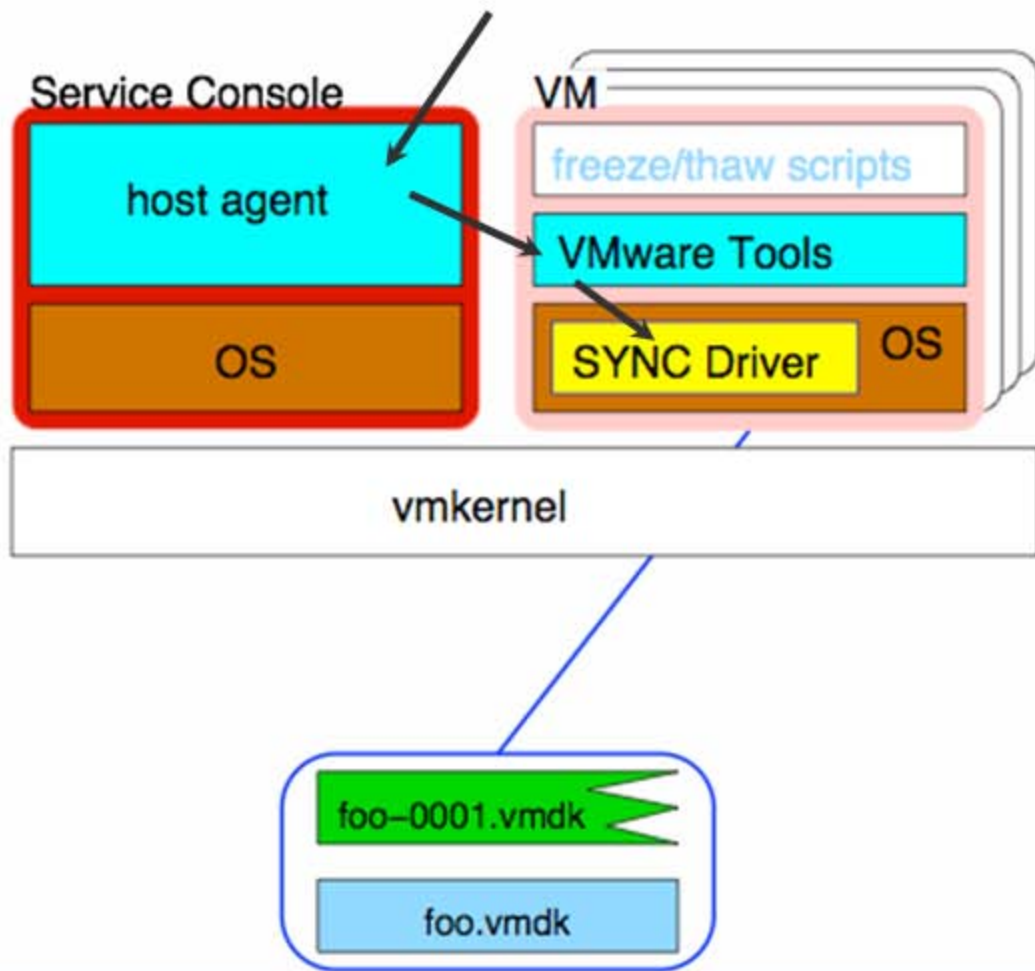
Snapshots in VI3

If there is a pre-freeze script, I'll run it

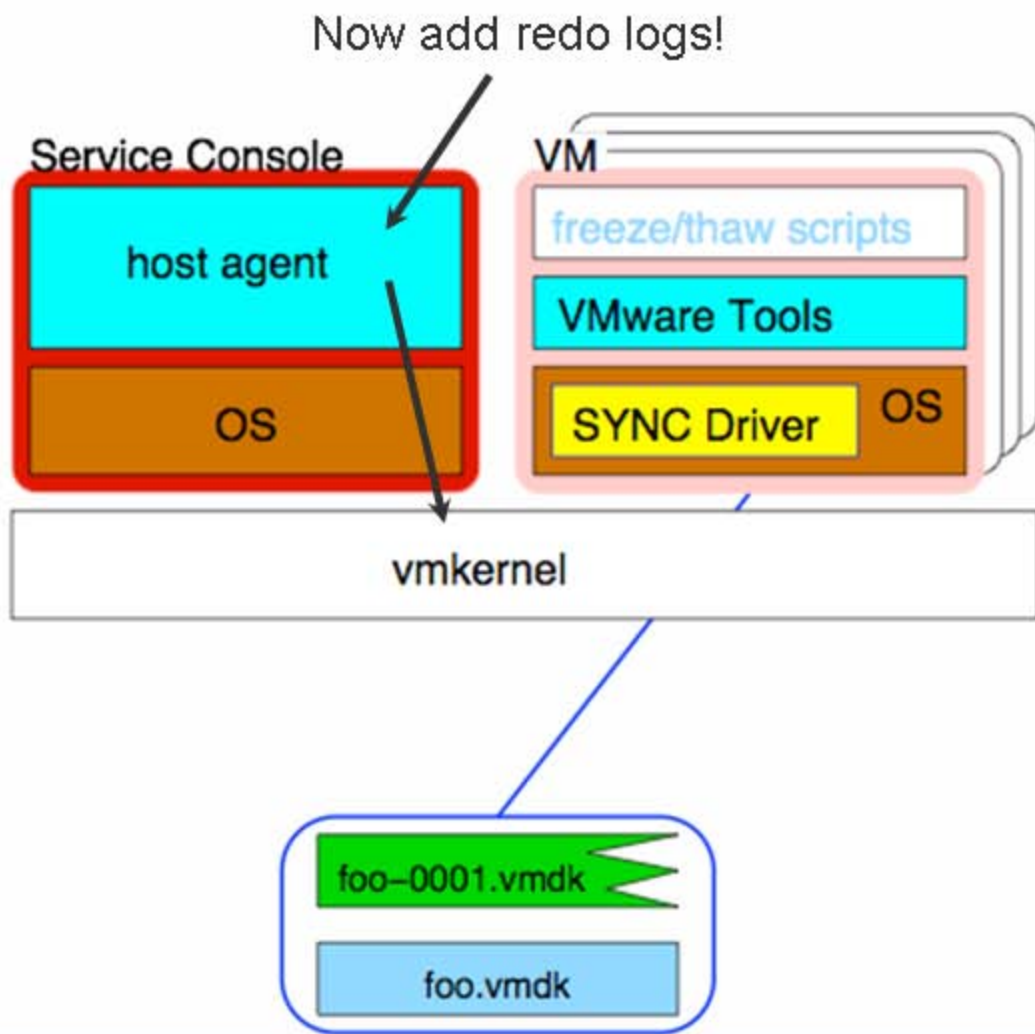


Snapshots in VI3

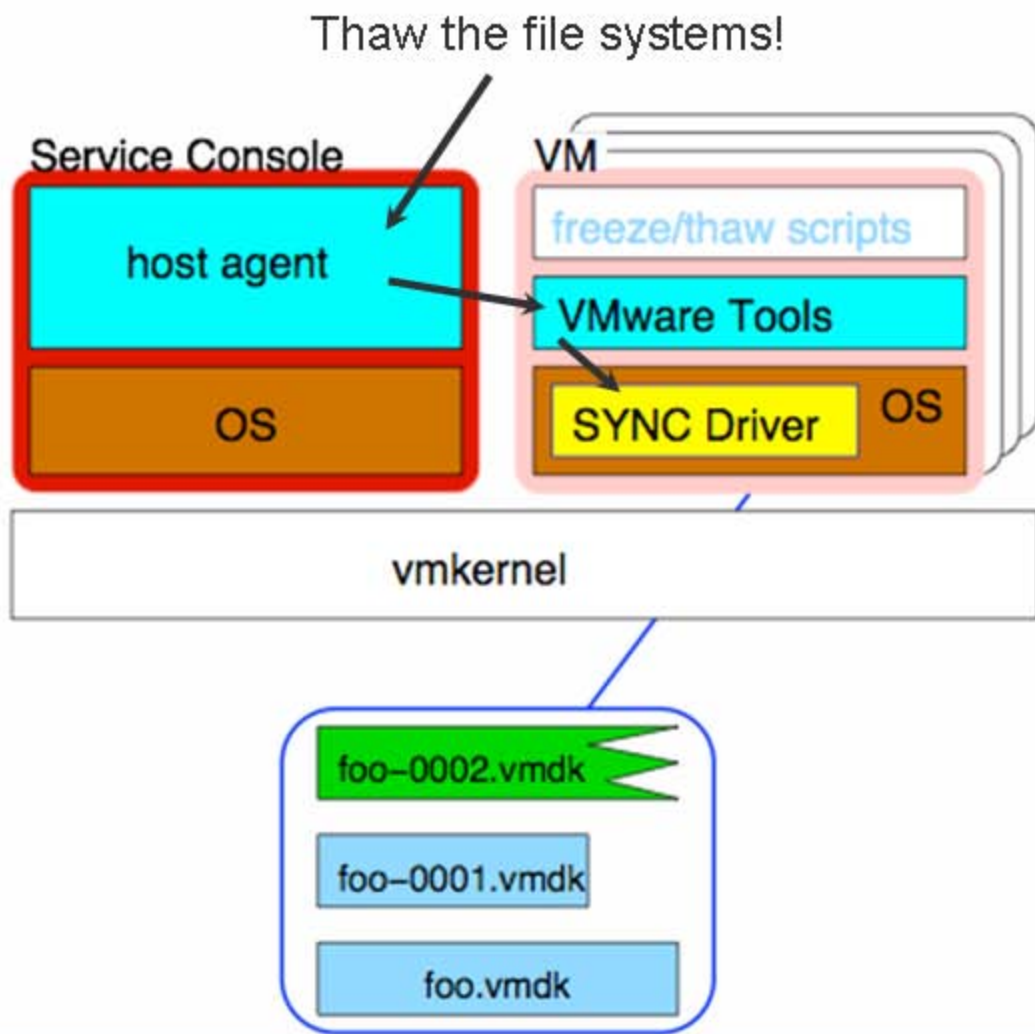
Time to quiesce the file systems!



Snapshots in VI3

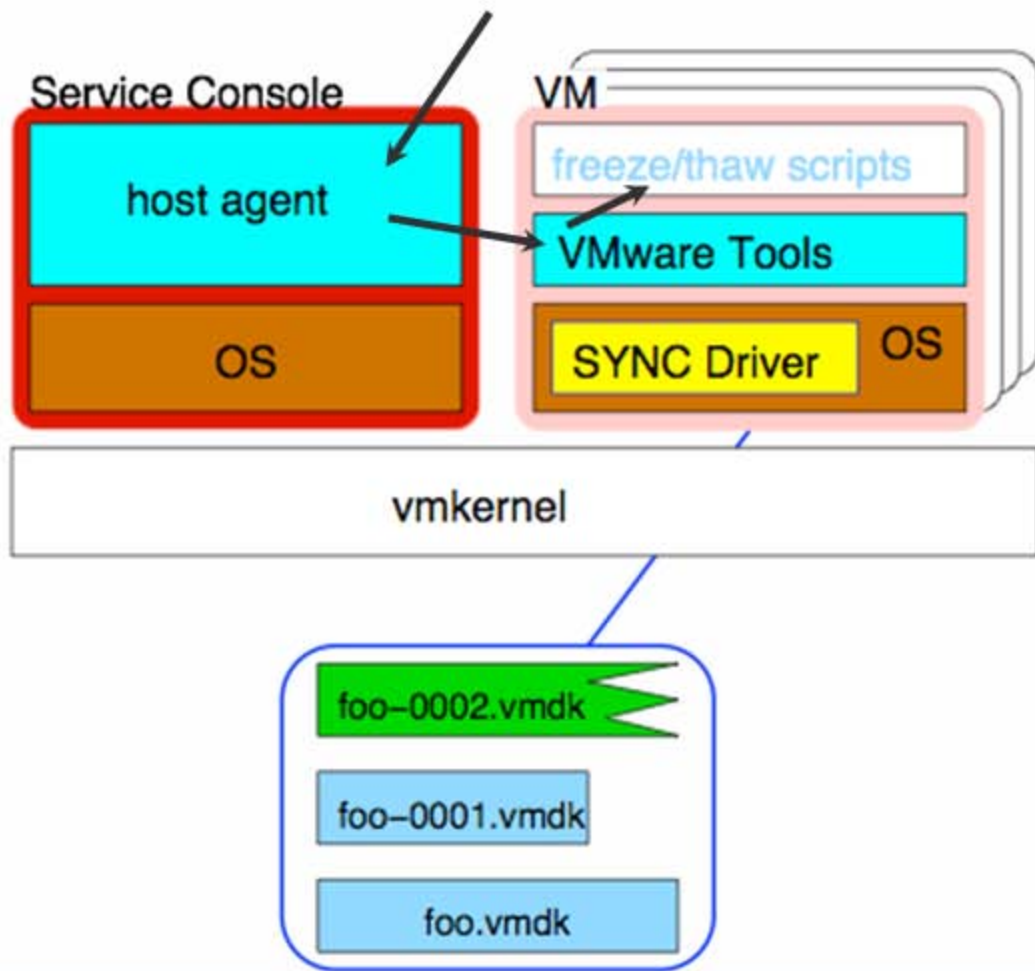


Snapshots in VI3



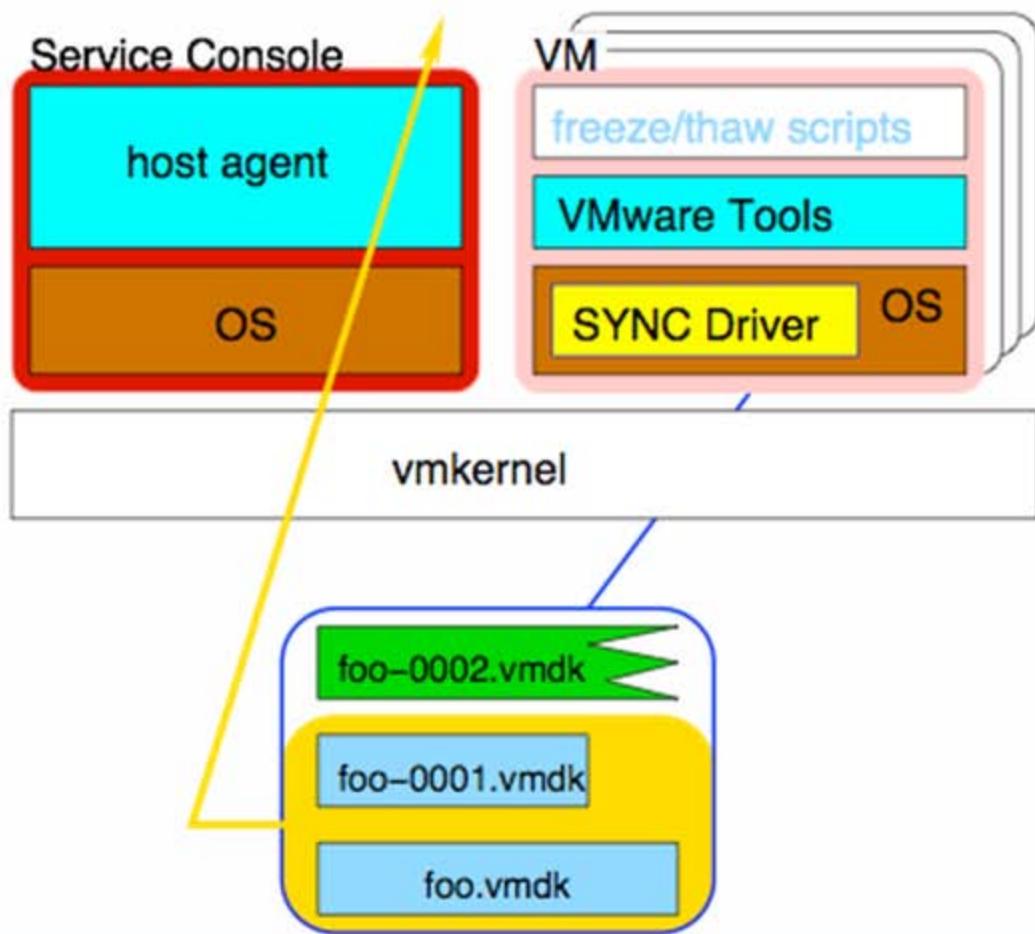
Snapshots in VI3

If there is a post-thaw script, I'll run it



Snapshots in VI3

There is your Snapshot!



Overview

- Classic vs. VM Aware Backup
- ESX Snapshots
- COW Files and COW Hierarchies
- Offloaded Backup with VCB
- What VCB is Not
- VCB and 3rd Party Backup Software
- What's Cooking in R&D

A Word About COW

- A virtual disk is represented by a VMFS base disk
- There can be COW files (redo logs) on top of a base disk ("COW hierarchy")
- Topmost redo log in COW hierarchy absorbs new writes
- All other COW files (+base disk) are immutable
- Building blocks for snapshots
- COW files are not continuous logs. Size can never exceed size of base disk

Overview

- Classic vs. VM Aware Backup
- ESX Snapshots
- COW Files and COW Hierarchies
- Offloaded Backup with VCB
- What VCB is Not
- VCB and 3rd Party Backup Software
- What's Cooking in R&D

Offloaded Backup with VCB

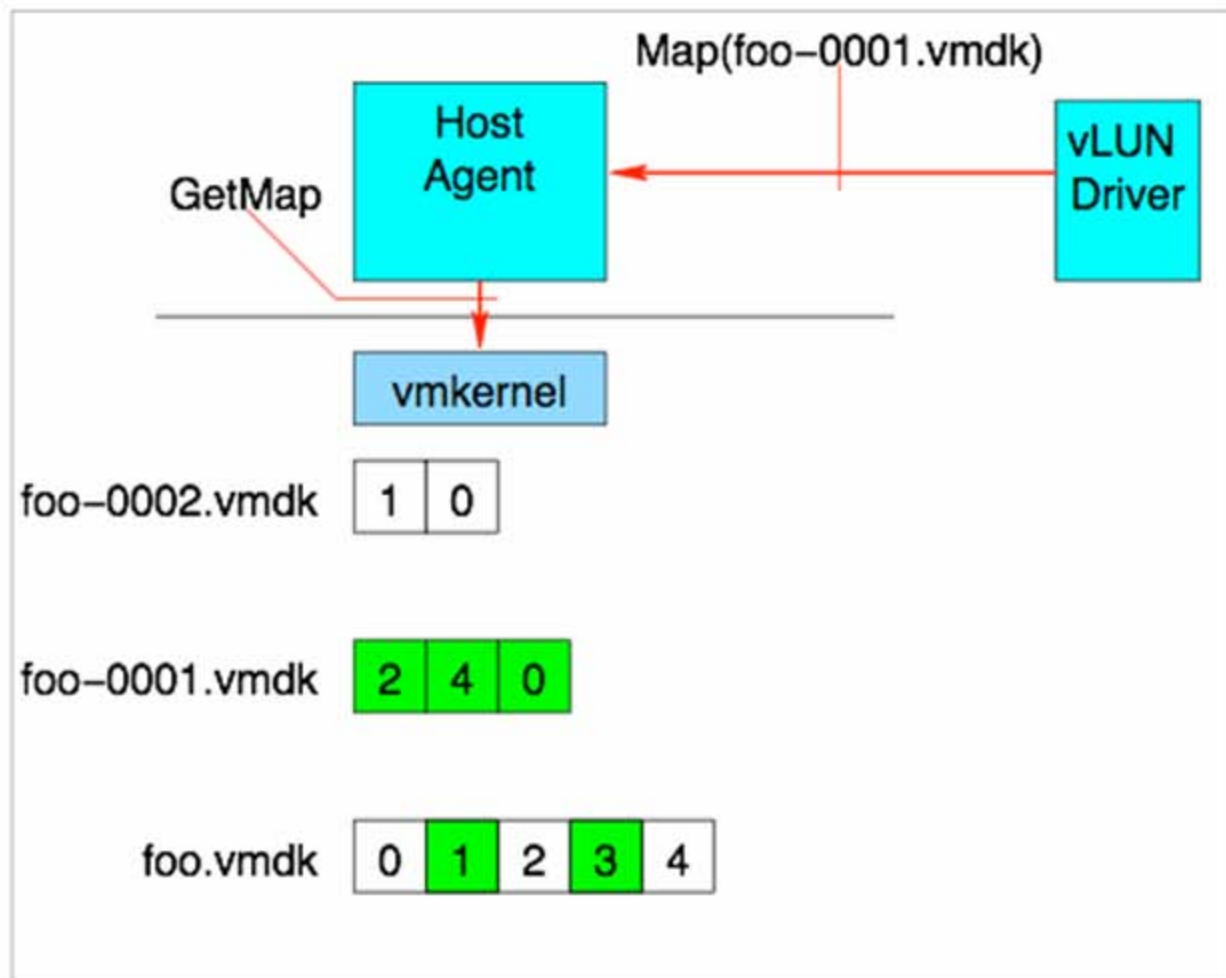
- VCB

- > Access VM snapshot from physical machine
- > Export disks in snapshot => full VM backup
- > Mount file systems on disks in snapshot (Windows) => file level backup
- > No load on ESX server(s), enables higher consolidation rates
- > Can be made entirely LAN-free

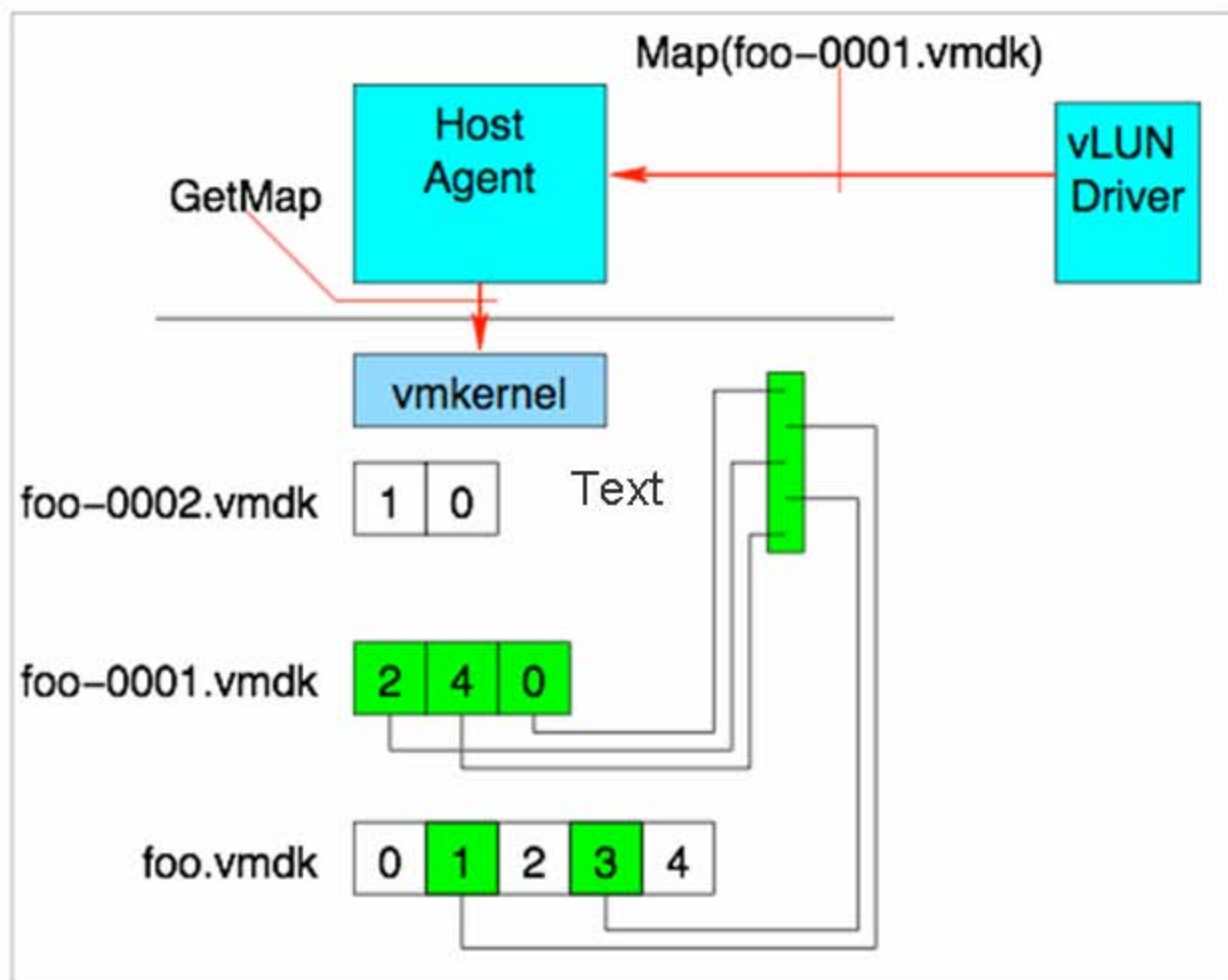
How Snapshot Data is Accessed

- There is no “VMFS Driver for Windows” on the proxy
- Data access uses metadata provided from ESX server hosting VM
 - > ESX Server provides translation from logical block number within virtual disk to SCSI LUN/offset on LUN (“vLUN driver”)
 - > VCB software on proxy uses this information to re-map reads to virtual disk
- In a server farm, VirtualCenter Server forwards request to appropriate ESX server
- VMotion is tracked by VirtualCenter Server

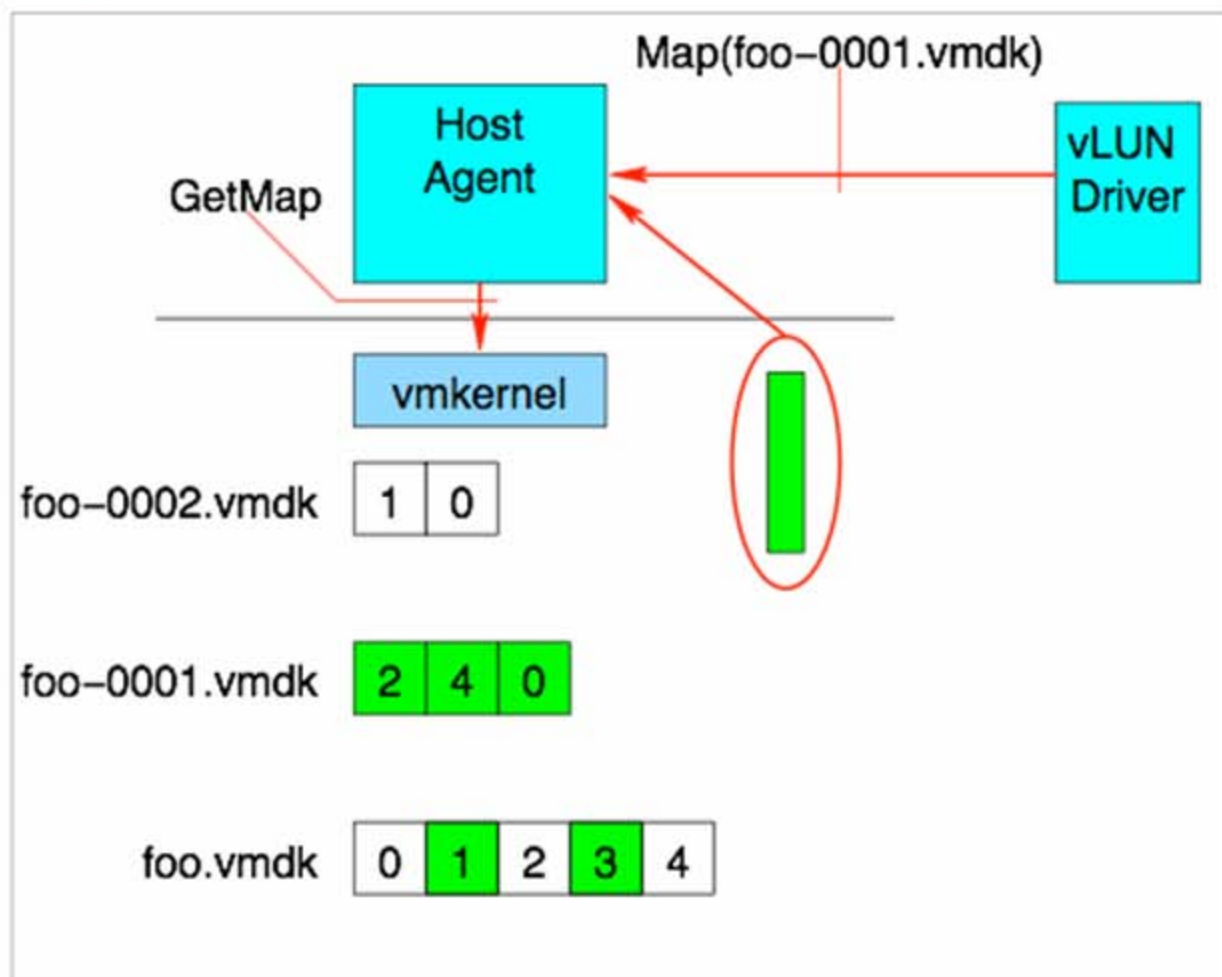
How Snapshot Data is Accessed



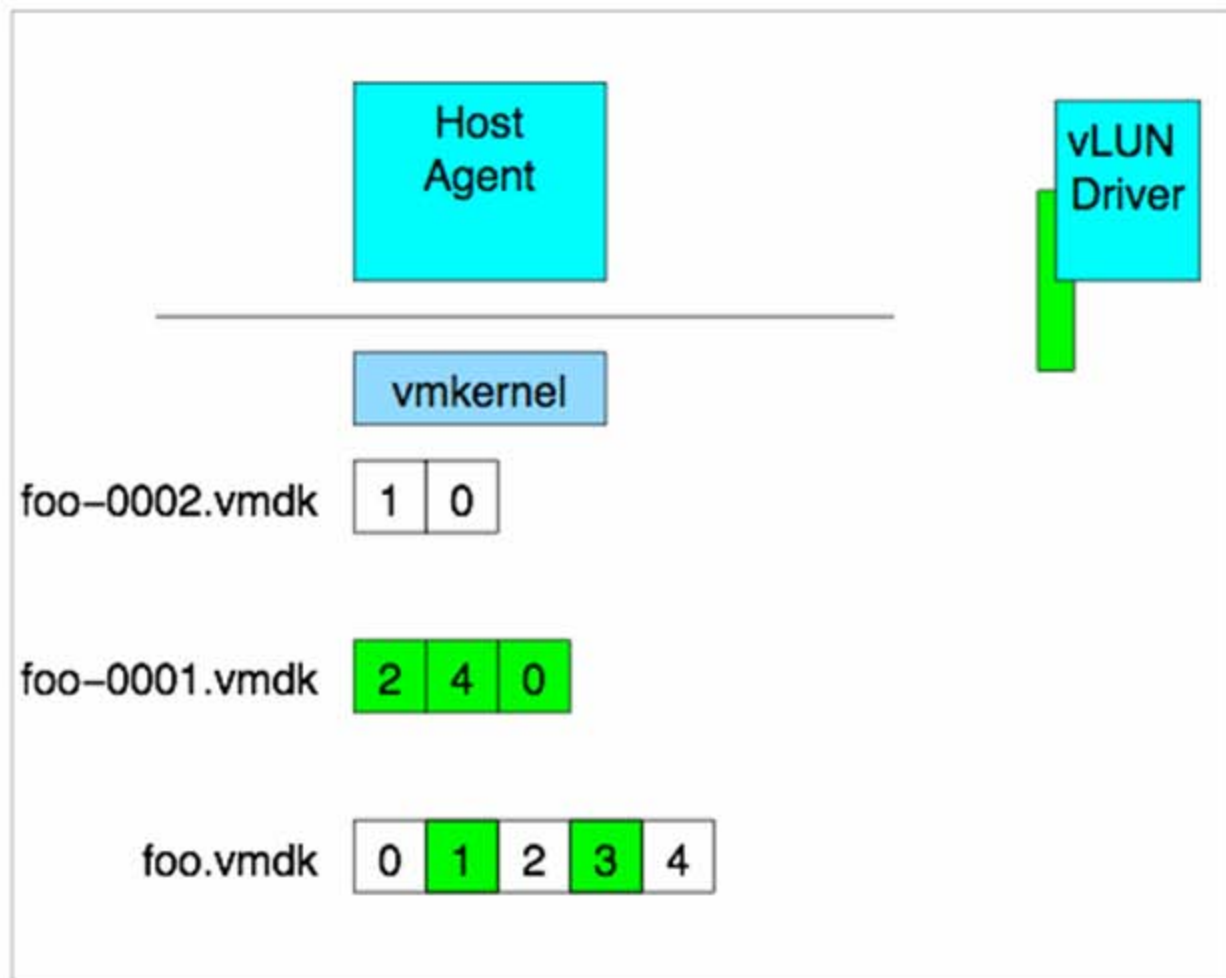
How Snapshot Data is Accessed



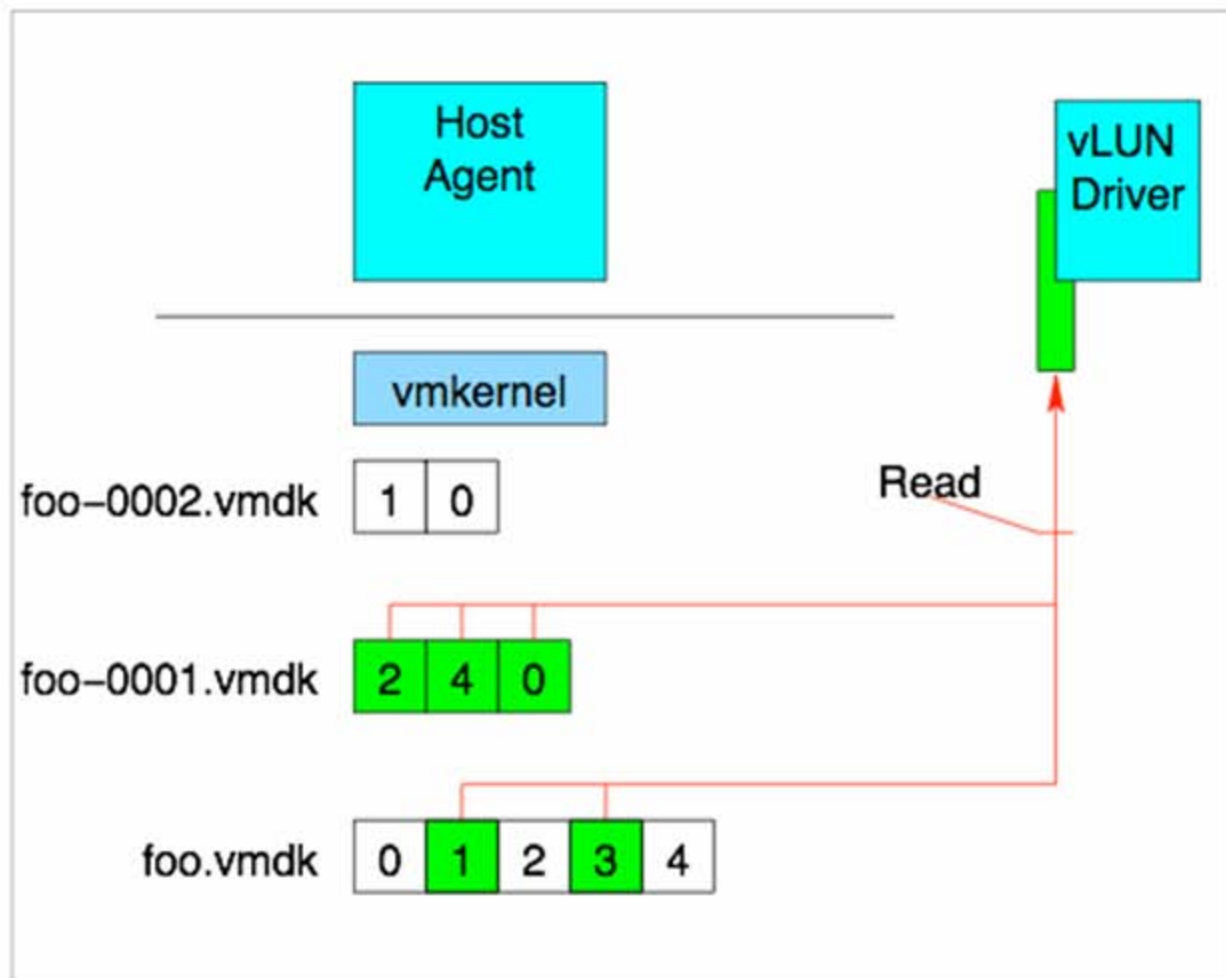
How Snapshot Data is Accessed



How Snapshot Data is Accessed



How Snapshot Data is Accessed



Overview

- Classic vs. VM Aware Backup
- ESX Snapshots
- COW Files and COW Hierarchies
- Offloaded Backup with VCB
- What VCB is Not
- VCB and 3rd Party Backup Software
- What's Cooking in R&D

What VCB is Not

- VCB is a highly efficient data mover
- Backup is more than data movement
 - > Catalog
 - > Scheduling
 - > Media/Backup Management
- VCB is not an out-of-the box backup solution
 - > Think of it as a toolkit
 - > Build your own solution (or have it built...)
 - > 3rd party integration modules enable end to end backup solutions

Overview

- Classic vs. VM Aware Backup
- ESX Snapshots
- COW Files and COW Hierarchies
- Offloaded Backup with VCB
- What VCB is Not
- VCB and 3rd Party Backup Software
- What's Cooking in R&D

VCB and 3rd Party Backup SW

- Backup agent running on Proxy uses VCB through pre and post backup scripts
- There is an Integration Module for each supported 3rd party backup application
- Integration Module ties into generic VCB framework
 - > JScript glue
 - > Command Line Utilities

VCB and 3rd Party Backup SW

- File Level backup (currently Windows only)
 - > Pre-backup script
 - Snapshots VM
 - Presents disk snapshots as “virtual LUNs” to OS on proxy
 - Controls mounting of file systems on proxy
 - Drive letter discovery
 - > Backup SW traverses mounted file systems, just as for any locally stored data

VCB and 3rd Party Backup SW

- File Level backup
 - > Post-backup script
 - Unmounts file systems
 - Deletes snapshot
 - > Challenges for Integration Module
 - Ownership of data (belongs to respective VM, not proxy) => put VM into groups; one alias for backup proxy per group
 - File level restore is at the realm of 3rd party backup SW: "redirected restore"

VCB and 3rd Party Backup SW

- Full VM backup (any guest OS)
 - > Pre-backup script
 - Snapshots VM(s)
 - Exports VM(s) onto proxy (WS-style disks, WS-like config info)
 - Deletes snapshot
 - > Backup SW traverses directories containing exported VMs
 - > Post-backup script
 - Deletes exported VM data

VCB and 3rd Party Backup SW

- Full VM backup (any guest OS)
 - > Restore currently through command line utility on Service Console

Overview

- Classic vs. VM Aware Backup
- ESX Snapshots
- COW Files and COW Hierarchies
- Offloaded Backup with VCB
- What VCB is Not
- What's Cooking in R&D

VCB - What's Cooking in R&D

- First, a disclaimer...
- These features are representative of feature areas under development. Feature commitments must not be included in contracts, purchase orders, or sales agreements of any kind. Technical feasibility and market demand will affect final delivery

VCB - Identifying LUNs

- At the mercy of the SCSI Layer...
- Cannot access disk: One LUN could not be opened
 - > Requirement for matching LUN IDs (ESX/Proxy) can make deployment difficult.
 - > SAN arrays can often be seen over multiple paths, but not all of them are usable for IO (active/passive)
 - > Storage virtualization (multipathing SW) can distort results of SCSI inquiries

VCB - Identifying LUNs

- Most VM Storage resides on VMFS
- VMFS metadata contains unique signatures
 - > Looking into using these signatures to identify LUNs
- Independent of SCSI inquiries
 - > Would remove requirement for matching LUN IDs
 - > No more issues with distorted SCSI inquiries
- Bonus: Attempt to read VMFS signatures allows detection of inactive paths
 - > Would deal with inactive paths more gracefully

VCB - What Else is Going on?

- Also under investigation
 - > How to make full VM restore faster/easier
 - > Can we do something to improve snapshot consistency even further?
 - Some applications should be quiesced
 - Quiescing of different guest OSes
- Many 3rd party backup SW vendors have integration modules or are working on them

Q&A

Presentation Download

Please remember to complete your
session evaluation form
and return it to the room monitors
as you exit the session

The presentation for this session can be downloaded at
<http://www.vmware.com/vmtn/vmworld/sessions/>

Enter the following to download (case-sensitive):

Username: cbv_rep
Password: cbvfor9v9r

Some or all of the features in this document may be representative of feature areas under development. Feature commitments must not be included in contracts, purchase orders, or sales agreements of any kind. Technical feasibility and market demand will affect final delivery.

VMWORLD 2006

