



Ensuring High Availability with VMware vSphere 4

David Marie

Systems Engineer

18 Mai 2009

Disclaimer

This session may contain product features that are currently under development.

This session/overview of the new technology represents no commitment from VMware to deliver these features in any generally available product.

Features are subject to change, and must not be included in contracts, purchase orders, or sales agreements of any kind.

Technical feasibility and market demand will affect final delivery.

Pricing and packaging for any new technologies or features discussed or presented have not been determined.

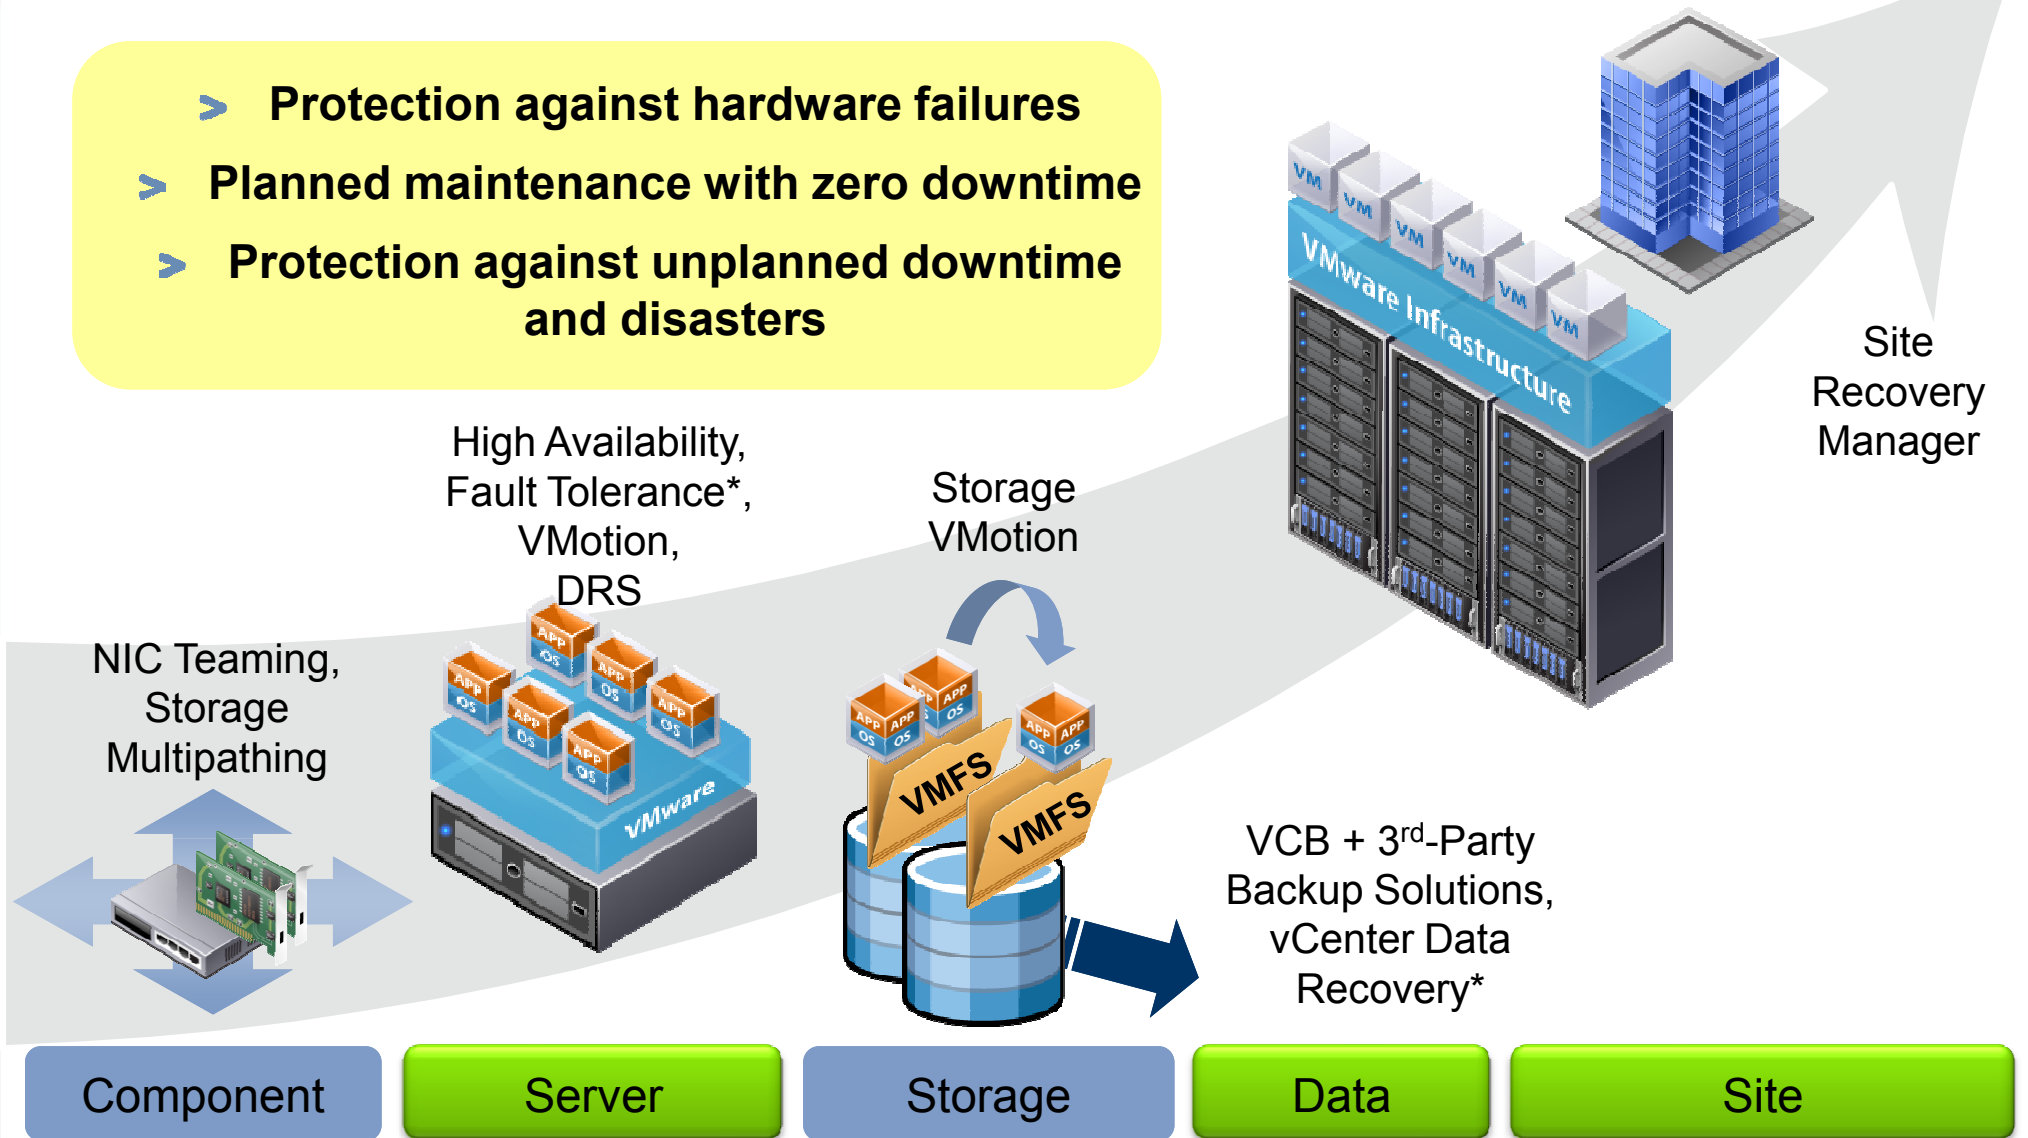
“These features are representative of feature areas under development. Feature commitments are subject to change, and must not be included in contracts, purchase orders, or sales agreements of any kind. Technical feasibility and market demand will affect final delivery.”

Outline

- High availability overview
- Basic concepts of virtual machine fault tolerance
- Hardware and software requirements
- FT Demo
- Performance, best practices and applications
- Q & A

VMware Offers Protection At Every Level

- > Protection against hardware failures
- > Planned maintenance with zero downtime
- > Protection against unplanned downtime and disasters



* Available in 2009

Planned vs. Unplanned downtime


Planned downtime through workload migration

- > VMware VMotion
- > VMware Storage VMotion
- > VMware Distributed Resource Scheduler (DRS)

Unplanned downtime with automated failover, backup & recovery

- > VMware Fault Tolerance (FT)
- > VMware High Availability (HA)
- > vCenter Data Recovery
- > VMware Consolidated Backup
- > VMware Site Recovery Manager
- > ESX availability features

VMware High Availability (VMware HA)



Datacenter Observations

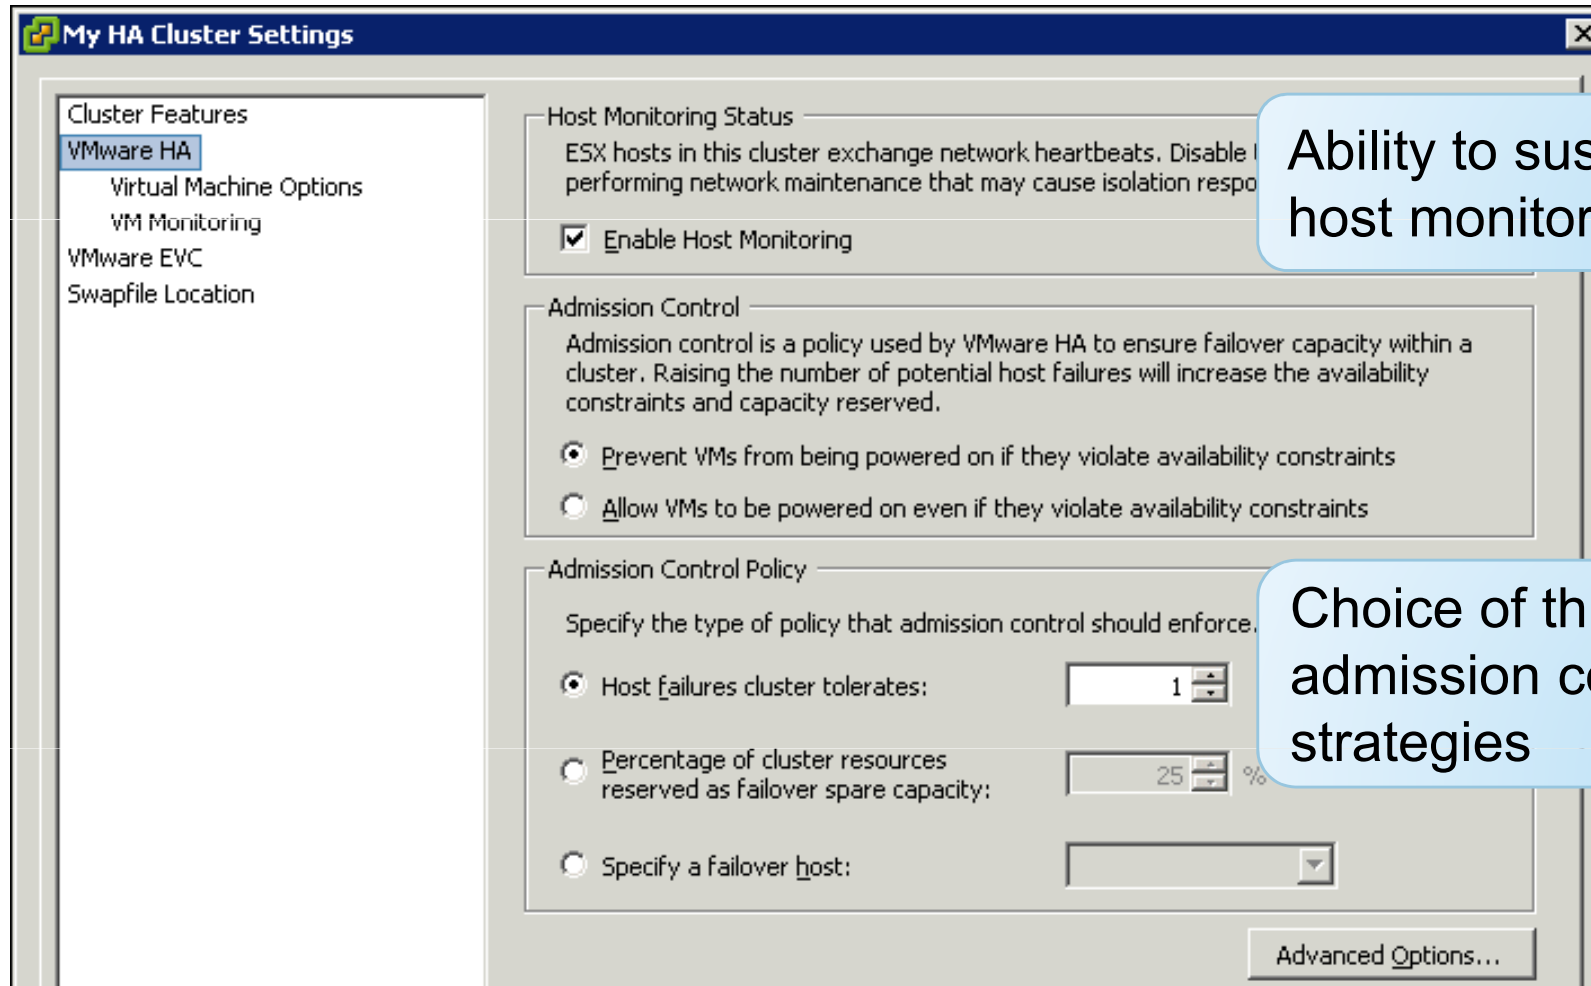
- > Hardware fails, humans make mistakes
- > Clustering is too complex / expensive
- > 90% of x86 workloads lack redundancy

Requirements

- > Minimize downtime and costs
- > Protect more workloads
- > Application & OS Independence
- > Simple deployment and management

VMware HA detects failures and provides rapid application recovery.

New for HA: Cluster Settings



Ability to suspend host monitoring

Choice of three admission control strategies

New for HA: VM Monitoring

The screenshot shows the 'My HA Cluster Settings' window. On the left, a navigation pane lists 'Cluster Features' with sub-items: 'VMware HA', 'Virtual Machine Options', 'VM Monitoring' (selected), 'VMware EVC', and 'Swapfile Location'. The main area is divided into sections: 'VM Monitoring Status' with a description and an 'Enable VM Monitoring' checkbox; 'Default Cluster Settings' with a 'Monitoring sensitivity' slider (Low to High) and a 'Custom' checkbox; a configuration box for 'Failure interval' (30 seconds), 'Minimum uptime' (120 seconds), 'Maximum per-VM resets' (3), and 'Maximum resets time window' (radio buttons for 'No window' and 'Within: 1 hours'); and 'Virtual Machine Settings' with a table of VMs and their monitoring settings.

Virtual Machine	VM Monitoring
W2K8-VM1	Custom...
W2K3DE-VM2	Use cluster settings
XPVM01	Use cluster settings
W2K3VM2	High
Thin Disk VM	Medium
	Low
	Disabled
	Custom...

Enable automatic restart due to failure of guest operating system


Determine how quickly failures are detected

Set monitoring sensitivity for individual virtual machines

Outline

- High availability overview
- Basic concepts of virtual machine fault tolerance
- Hardware and software requirements
- FT Demo
- Performance, best practices and applications
- Q & A

VMware Fault Tolerance (FT)




Exciting technology for mission critical VM's

- Zero downtime and zero data loss in case of hardware failures
- Single VM running on two hosts for redundancy
- Redundancy automatically restored after failover
- Simple to deploy

First release:

- All ESX supported Guest Operating Systems
- Single vCPU workloads
- Integrated with VMware HA
- Multiple FT VMs per host
- Dynamically enabled or disabled per VM

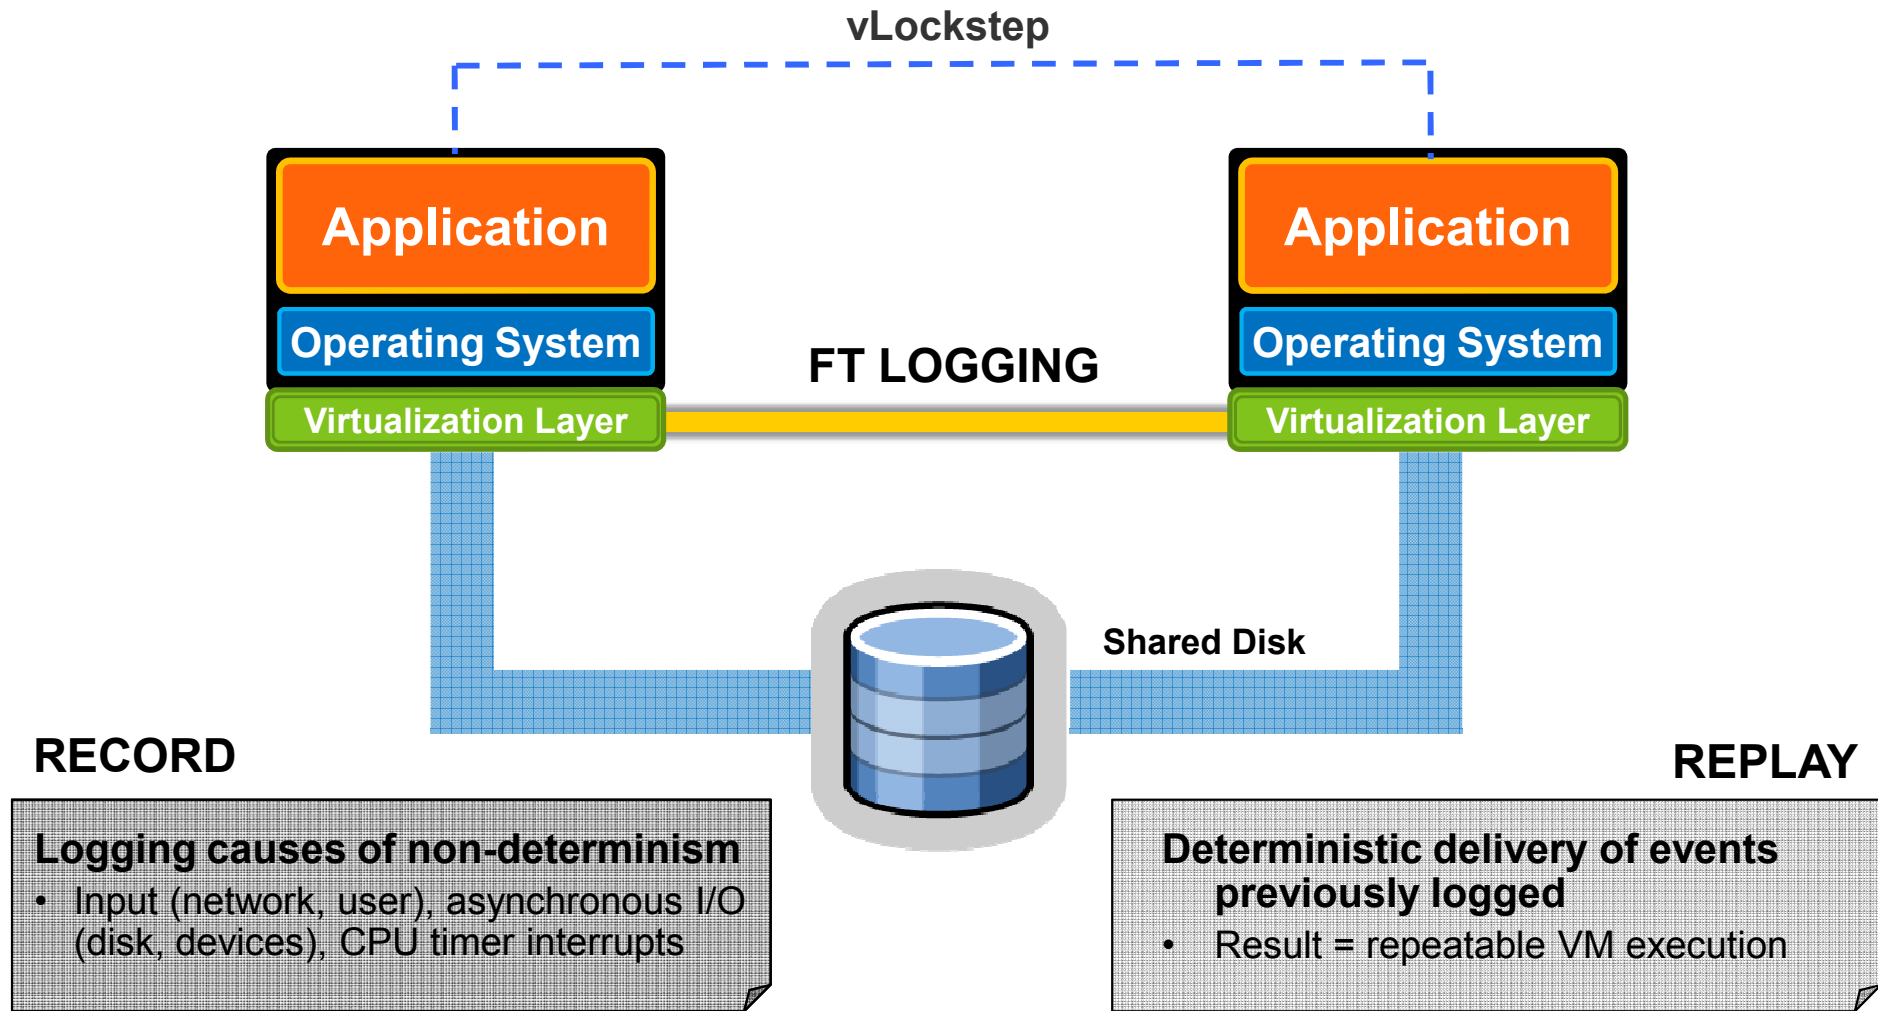
VMware Fault Tolerance (FT)



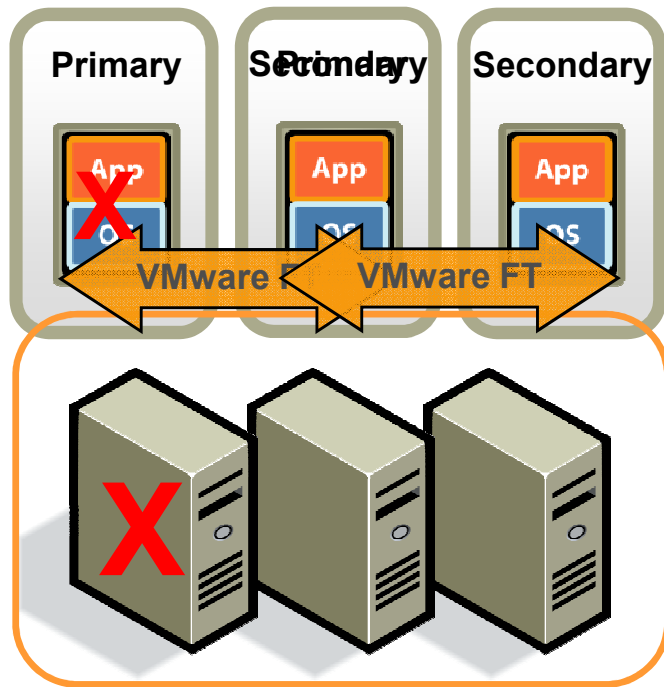
Technology Overview

- > VMware FT protects a VM (Primary) by creating a copy on another node (Secondary)
- > If the Primary's server fails, the Secondary VM takes over without downtime or loss of data
- > VMware FT automatically creates a new Secondary to continue protecting the VM
- > vLockstep technology guarantees that Primary and Secondary are in full synchronization

vLockstep Technology

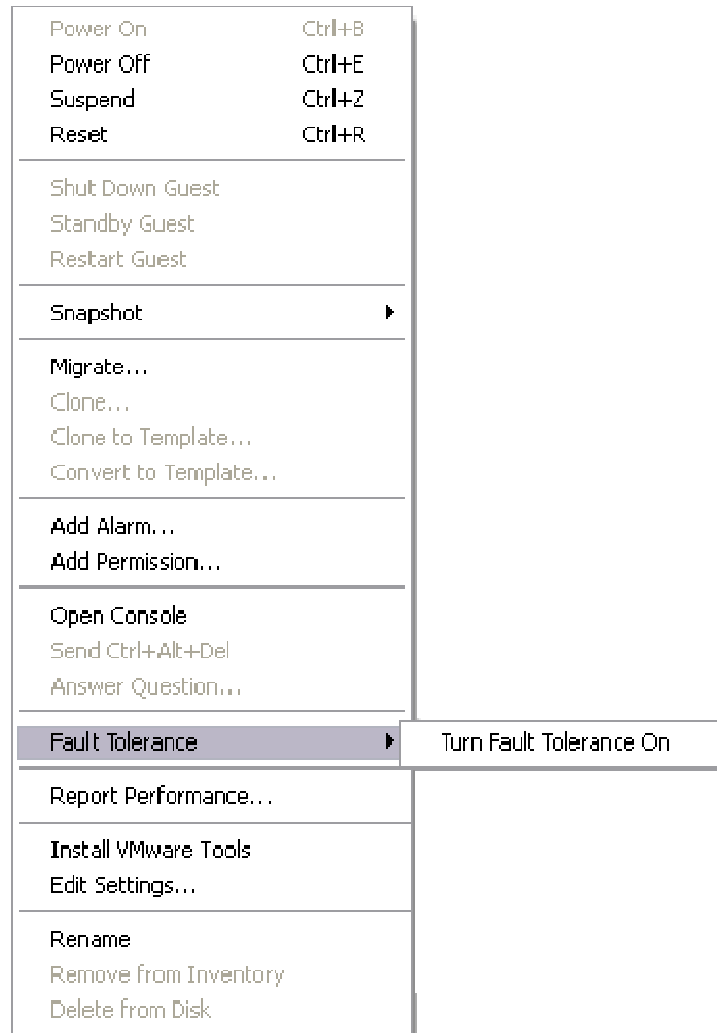


VMware FT Operation



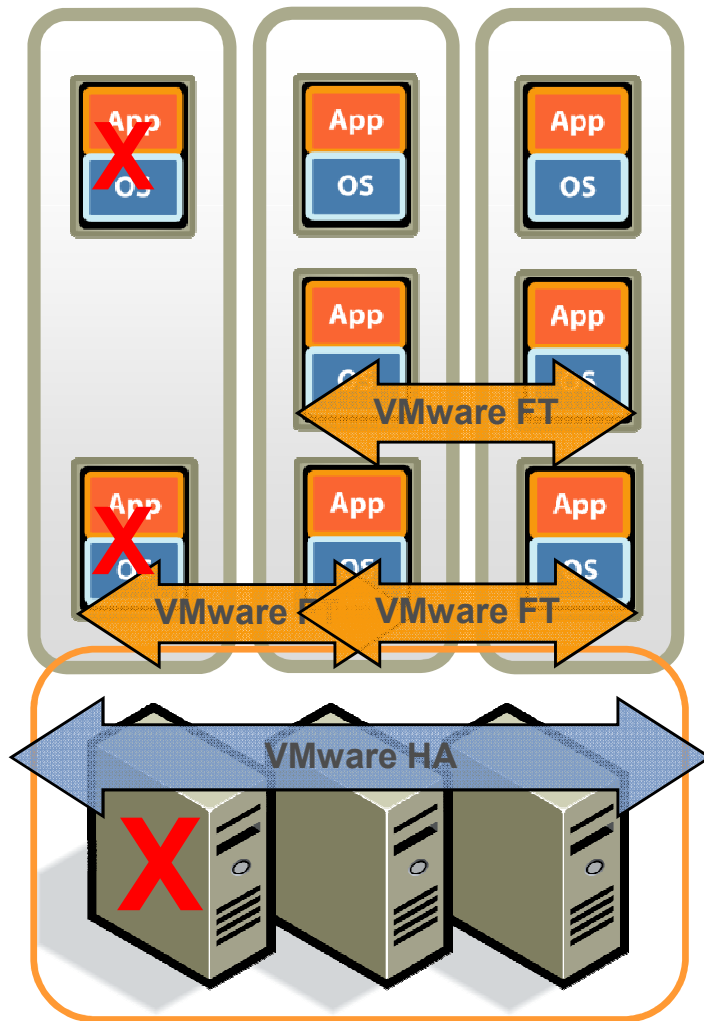
- > Primary and secondary continuously monitor each other. Failures are noticed within a few seconds
- > Secondary becomes the new “primary”. Shared disk operations guarantee reliable takeover (no “split brain”)
- > Network behavior is like VMotion ... no loss of connection, no lost data

Simple Configuration



- > Enabling FT is simple and straightforward
- > Simply select the VM and select FT Enable VMware FT does the rest
- > The Secondary VM is automatically created on an eligible node and you are instantly protected

VMware FT and HA Work Together




- > VMware HA protects all nodes of the cluster
- > Mission-critical VMs are protected by FT, remaining VM's protected by HA

When a host fails:

- > FT secondary takes over
- > New FT secondary is started by HA
- > HA-only VM's are restarted

Flexible Fault Tolerance

- Primary and secondary can be migrated using vMotion to other hosts
- Fault tolerance is managed like a VM property. It can be enabled/disabled at any time
- Storage vMotion requires FT to be disabled



Outline

- High availability overview
- Basic concepts of virtual machine fault tolerance
- Hardware and software requirements
- FT Demo
- Performance, best practices and applications
- Q & A

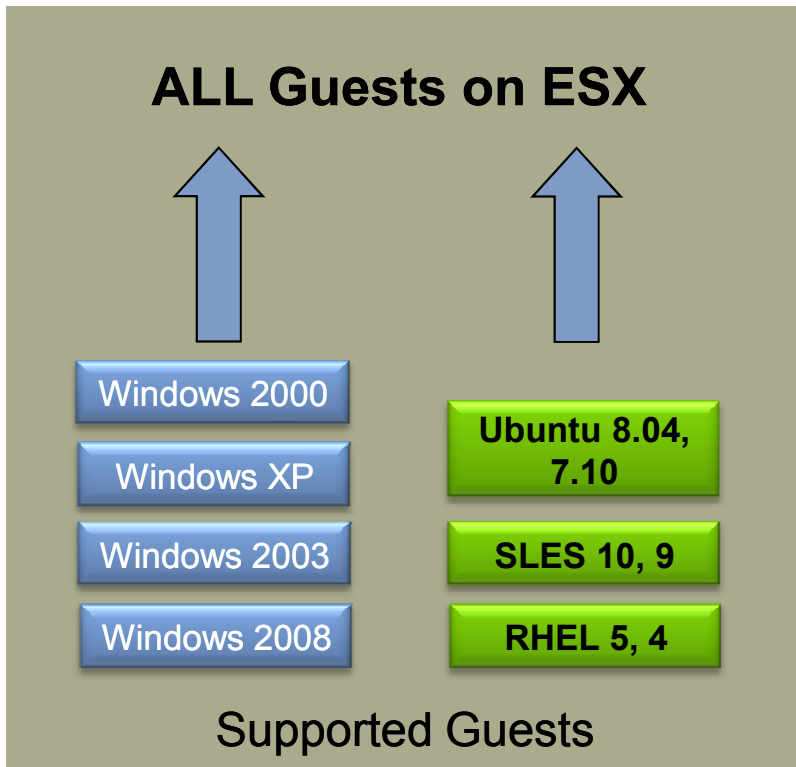
Host Requirements

Feature	Requirement
CPU	<ul style="list-style-type: none">• AMD Barcelona or• Intel Penryn• HV-enabled
Storage	Shared FC, iSCSI, or NAS
NIC	2 FT Logging NIC (suggested) 1 Gb or better
ESX	Same build on each host
Features	VMware HA

>50% of servers on
HCL support FT

55% of customers use
HA in production

Virtual Machine Requirements



- > **NO special Guest drivers**
- > **NO special Guest patches**

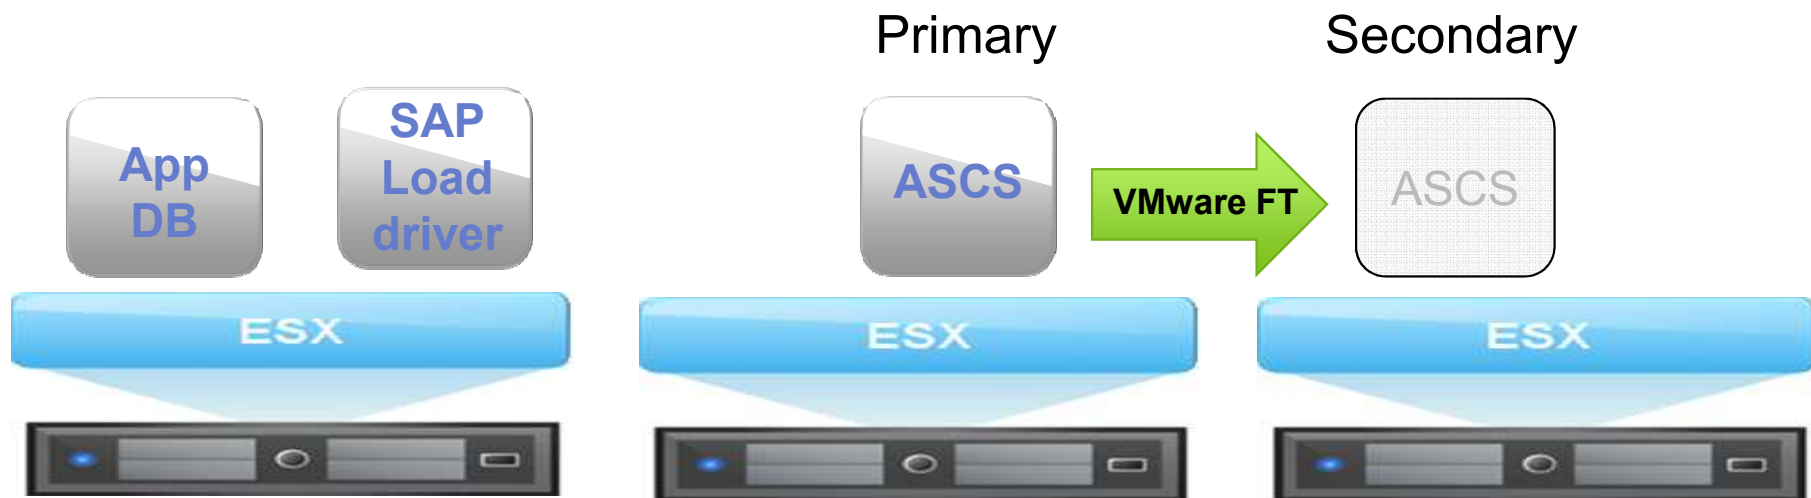
Feature	Requirement
Guest OS	ALL ESX supported 32-bit and 64-bit
vCPU	1
Virtual Disks	“Thick”, i.e. not thin-provisioned
Paravirt	Not supported
Unsupported Devices	USB Sound CD-ROM (physical) Floppy (physical) RDM (physical)
VMware Tools	Recommended

Outline

- High availability overview
- Basic concepts of virtual machine fault tolerance
- Hardware and software requirements
- FT Demo
- Performance, best practices and applications
- Q & A

SAP and VMware FT

- > SAP ECC 6.0 System based on SAP NetWeaver 7.0 platform
- > Single point of failure
 - ASCS = Message and Transaction locking service




Demo

Outline

- High availability overview
- Basic concepts of virtual machine fault tolerance
- Hardware and software requirements
- FT Demo
- Performance, best practices and applications
- Q & A

Preliminary FT Performance Results



Best Practices for Performance

- > Bandwidth of logging NIC could become a bottleneck
 - Logging NIC requires roughly [50 KB/sec + incoming network traffic] per FT VM
 - Don't put too many FT VMs on the same host (mix FT and non-FT VM's)
 - Mix primary and secondary VMs on same host
 - Consider 10 gigabit NIC for logging NIC if available
- > Lower round-trip latency on logging NIC improves guest networking performance


Target VMware FT Applications

Workload Type	Application	Rationale
Database	Small to medium instances that are strategic to IT infrastructure	Costs to deploy traditional cluster solutions not justified but availability is a must
Exchange and messaging	< 1000 users	Reduced licensing and management costs
Remote Branch Office	Many workloads	SLA requirements require a traditional cluster (\$\$\$). Deliver high availability at lower cost and easier to administer.
Custom applications	Business-specific solutions	Cluster solutions not available today

FT On-Demand



- > A VM can be FT-protected at any time
- > VMware FT automatically creates a Secondary VM and within a few minutes the VM is fully protected
- > Protection can be likewise disabled. When a FT for a VM is disabled, the Secondary VM is de-provisioned and resources are returned



Blackberry and VMware FT

- > Blackberry Enterprise Server 4.1.6 (BES)
- > Single point of failure
 - Forwards messages to Blackberry devices



Conclusion

VMware vSphere 4 increases the availability of your applications

- Storage multipathing and NIC teaming, VMotion, VMware HA, VMware Consolidated Backup, VMware Site Recovery Manager

VMware Fault Tolerance adds another level of availability for your datacenter

- Fault-tolerant VMs survive a hardware failure with no downtime or data loss
- Simple configuration and operation—easy to protect a VM with FT as needed
- Integrates with other features of vSphere 4 for maximum flexibility

Q&A