Veritas Cluster Server 6.0

New Features and Capabilities

Anthony Herr

Product Manager - Availability Products

What does function does VCS perform?

- High Availability
  - Ensure an application, in either a physical or virtual environment maintains its availability
  - Supports an application's Service Level Agreement (SLA)
- Application Management
  - Manage applications without in-depth application knowledge
  - Ensure an application stays online even without a failover target
- Single Pane-of-Glass Configuration and Control
  - Allow an enterprise view and manageability across UNIX/Linux/Windows hosts
  - Enables notification, health checks and reporting over the enterprise
- Multi-Tier Application Support
  - Enable top to bottom visibility of an application that crosses OS boundaries
  - Provide cross application dependencies on different cluster

Before we start, have you heard about…

How does VCS Monitor Applications?

- Legacy(Poll Based) Monitoring
  - Checks to see if the application is online or offline over an interval of time
  - Attributes for VCS monitoring which are controlled per resource type:
    - MonitorInterval when the application is online (default 60 sec)
    - OfflineMonitorInterval when the application is offline (default 300 sec)
    - MonitorTimeout is the amount of time given to a monitor process before giving up (default 60 sec)
  - Resources are monitored on all systems they are configured to run on
    - If an Oracle database is configured to run on a 3 node cluster then each of the three systems will validate the state of the resource based on the current resource state—online/offline
  - Each instance of a resource is monitored
    - If there are 20 mount resources in a service group, then 20 monitors will be run per system in the cluster based on the current resource state—online/offline

From polling to asynchronous monitoring

Faster failure detection

Traditional Monitoring Framework

Polling

Most Clustering Solutions

Poll based Monitoring

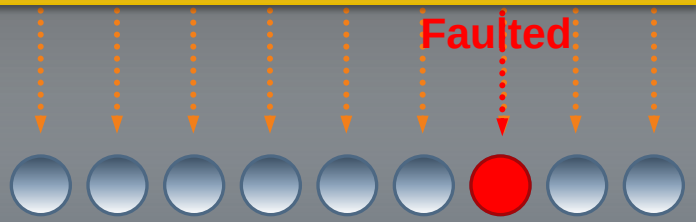

Resources Being Monitored

Intelligent Monitoring Framework

Asynchronous

Veritas Cluster Server (5.1 SP1 & 6.0)

Intelligent Monitoring

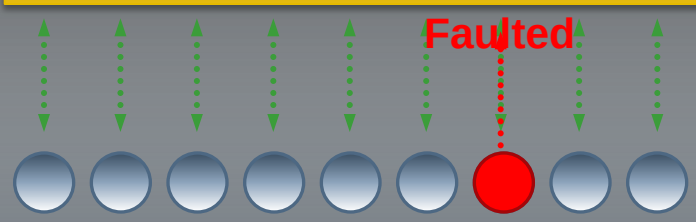

Resources Being Monitored

- Immediate fault detection
- Zero polling overhead
- Single step enablement

IMF Enabled Agents for UNIX/Linux:

- VCS 5.1 SP1
  - Process based agents
    - Physical environments, containers
      - IMF is enabled for Process agents running within a container
    - Oracle agent, Netlsnr agent
    - CVMvxconfigd
    - DB2 agent with VCS 5.1SP1RP1
  - Mount based agents
    - Mount, CFSMount
  - Application agent
    - Using PidFiles or MonitorProcess for Application Agent monitoring

IMF support is continually expanding as updates to agents are introduced in quarterly agent packs

- VCS 6.0
  - Virtualization based agents
    - Solaris Zones
    - AIX WPAR
  - Application agents
    - Sybase
  - IMF updates
    - Agent Framework update for Custom Agent support
    - **Support for IMF-PCV**

    Prevention of Concurrency Violation

    Application Agent with MonitorProcesses

- VCS 6.0.1
  - OpenIMF
  - DiskGroup Agent

IMF coverage for Windows

- VCS 6.0
  - GenericSerivce
  - ServiceMonitor
  - IP
  - NIC
  - MountV
  - Mount
  - VMDg
  - Oracle
  - NetLsnr
  - Process
  - RegRep
  - SQLServer2005
  - SQLAgService2005
  - SQLOlapService2005
  - MSDTC
  - SQLServer2008
  - IIS
  - ExchService2007
  - Exchange2010DB

Maximize Availability

Prevent Concurrencies (IMF-PCV)

Node with Application Online

VCS brings resource online

Agent registers with IMF

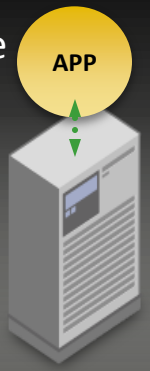

Registering...

Application is online

- **Intelligent Monitoring Framework – Prevention of Concurrency Violation (IMF-PCV)**
- **Prior to IMF-PCV, concurrency detection took up to 5 minutes**
- **No additional monitoring is required as this functionality is built on top of the IMF framework**

Node with Application Offline

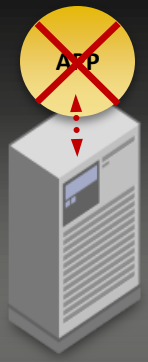

User attempts to bring app online outside of cluster

IMF recognizes attempt

PCV prevents

Application is Offline

- **Application Agent currently supported**
- **VCS is the only clustering technology with this feature**
- **In 6.0.1, both the Start Script and the Monitored Process are prevented**

New for 6.0

Enhanced Service Group Dependencies

Extend VCS to meet your specific application needs

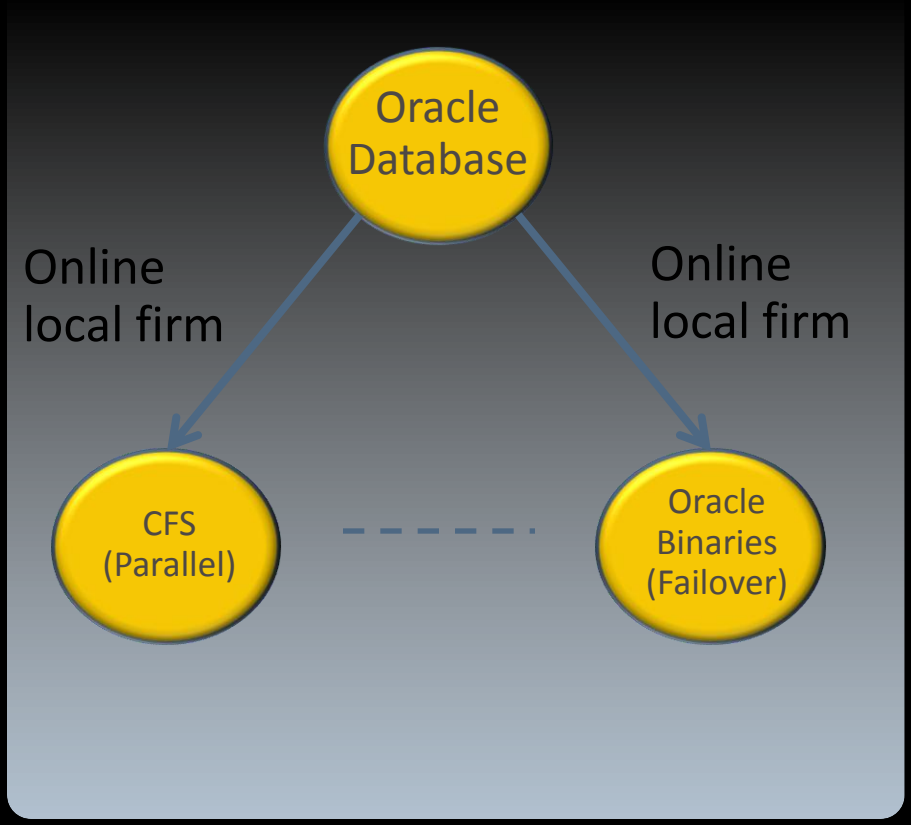

- Parent supports multiple child service groups
- Various combinations of dependencies (soft, firm) supported

Improvements to Trigger Events

Extend VCS to meet specific application needs

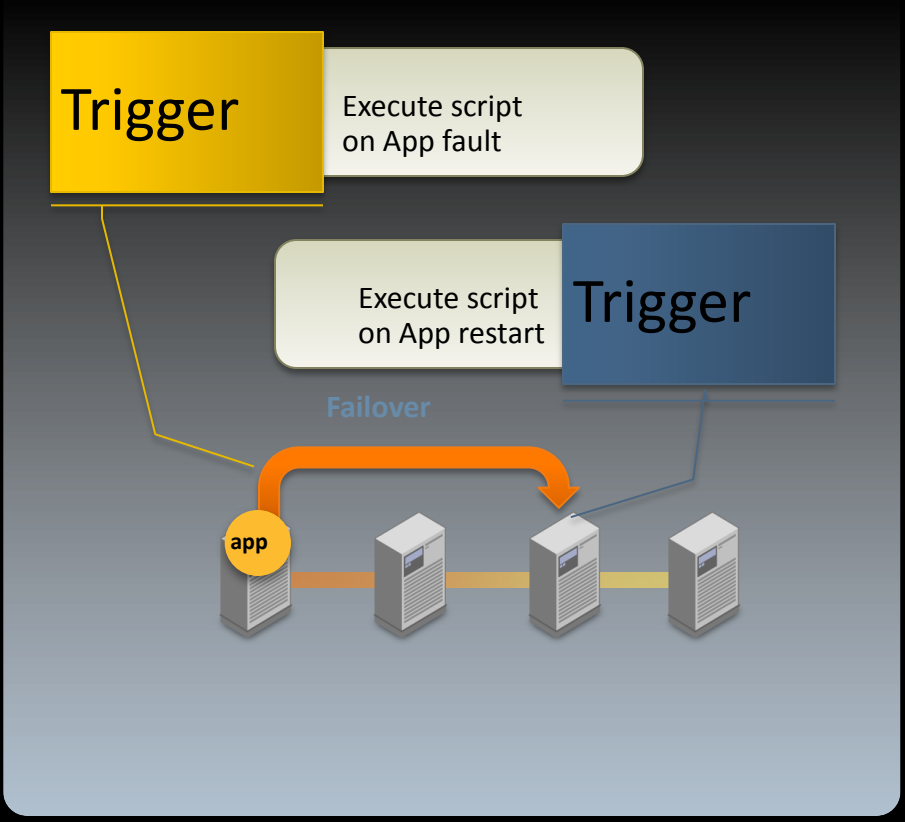

- Execute custom scripts for HA events
- New trigger for Resource Restarted event
- New support for multiple scripts to be executed in specified order

I/O Fencing

Fencing Methodology Comparison

Fencing Technology	Advantages	Notes
SCSI3 based fencing	• Data Protection is Fool-Proof in using of SCSI3-PR keys • Both SCSI3 Disks and CPS can be used for split-brain protection	• Requires SCSI3-PR capable disks • Not all virtual environments support SCSI3-PR
Non-SCSI3 Fencing (NSF)	• Supported in all virtual and physical environments • CPS (Coordination Point Server) is used for cluster membership arbitration • Each CPS can server multiple clusters reducing the number of disks needed	• Judicious use of timing to provide data protection • Provides network based membership arbitration • SCSI3-PR not used

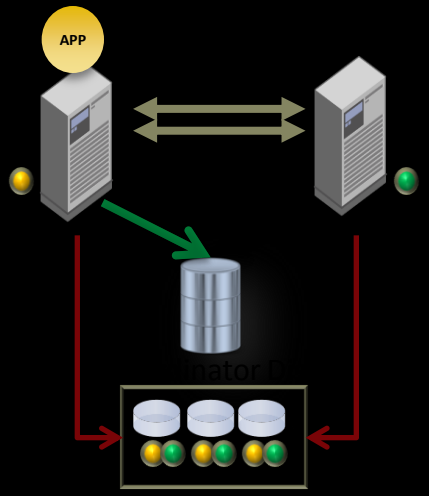

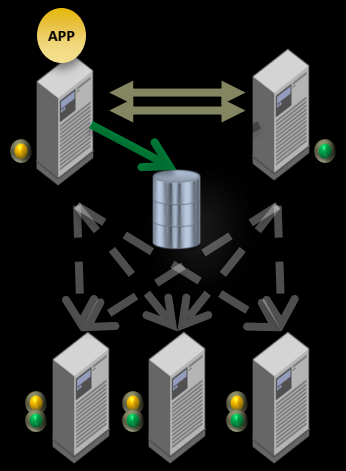

From disk based to network based arbitration

Reduced data risk in more environments

SCSI3 Disk Based Arbitration

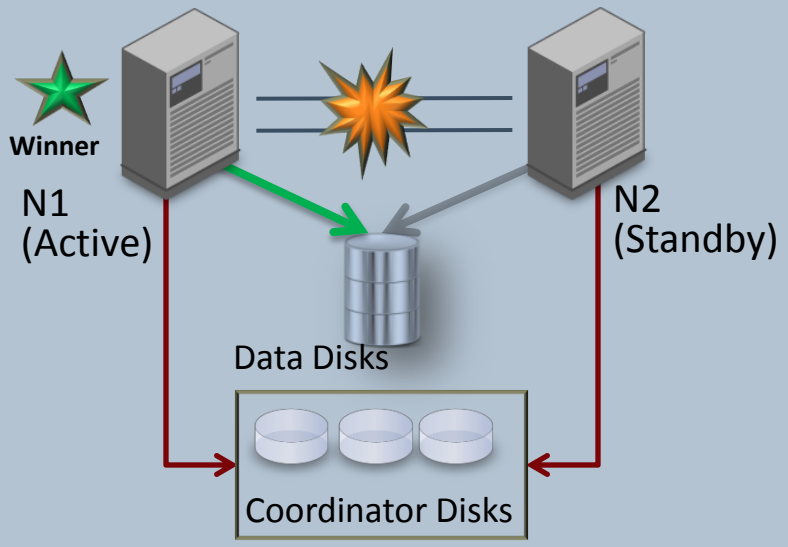

- SCSI3 protocol enables data protection
- Losing node prevents access to disk and then panics to prevent splitbrain

Coordination Point Server

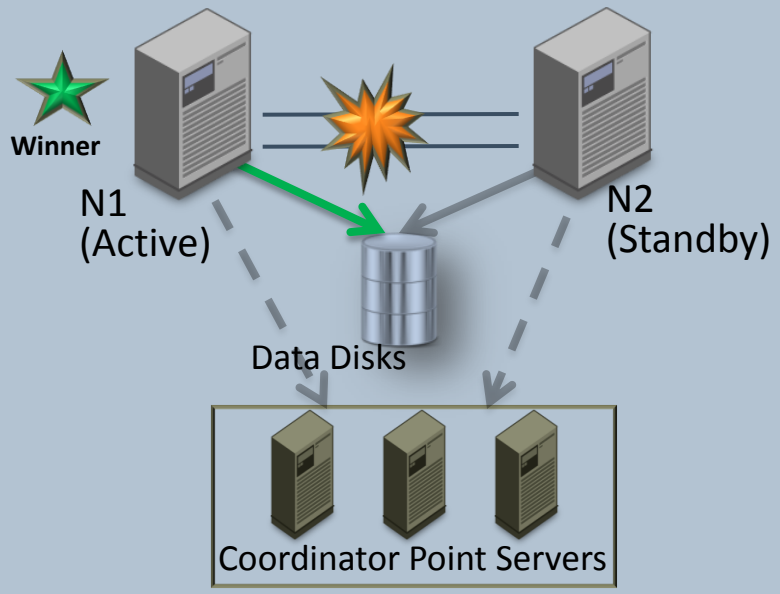

- Same CP Servers for many clusters
- SCSI3 + CPS arbitration also supported
- Virtual environments, campus clusters

I/O Fencing: Membership Arbitration & Data Protection

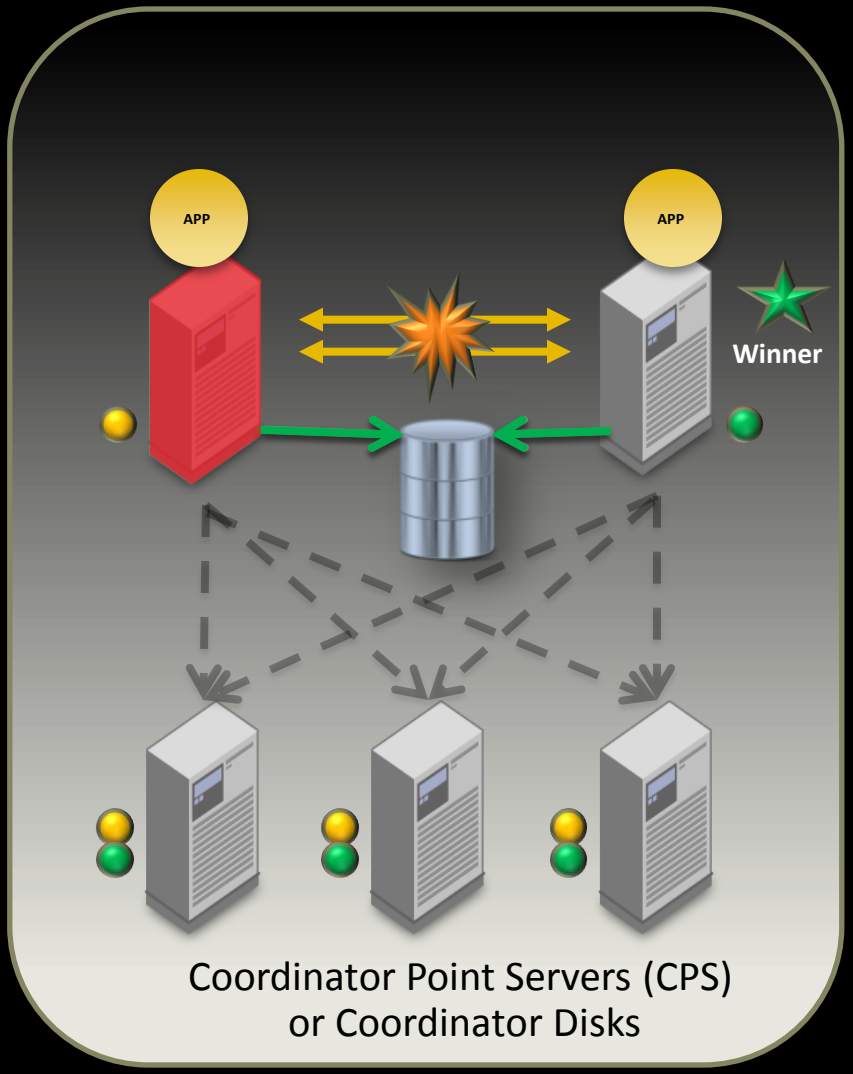

Split-Brain Protection

- On membership change, each subcluster will elect a race node to race for the coordination points
- The first racer node to reach the coordination point will remove the registration of other racer node
- The racer node which has its registration on greater than ½ of the configured coordination points is winner
- The losing racer node alerts all subcluster nodes to panic, resolving the network partition

I/O Fencing: Membership Arbitration

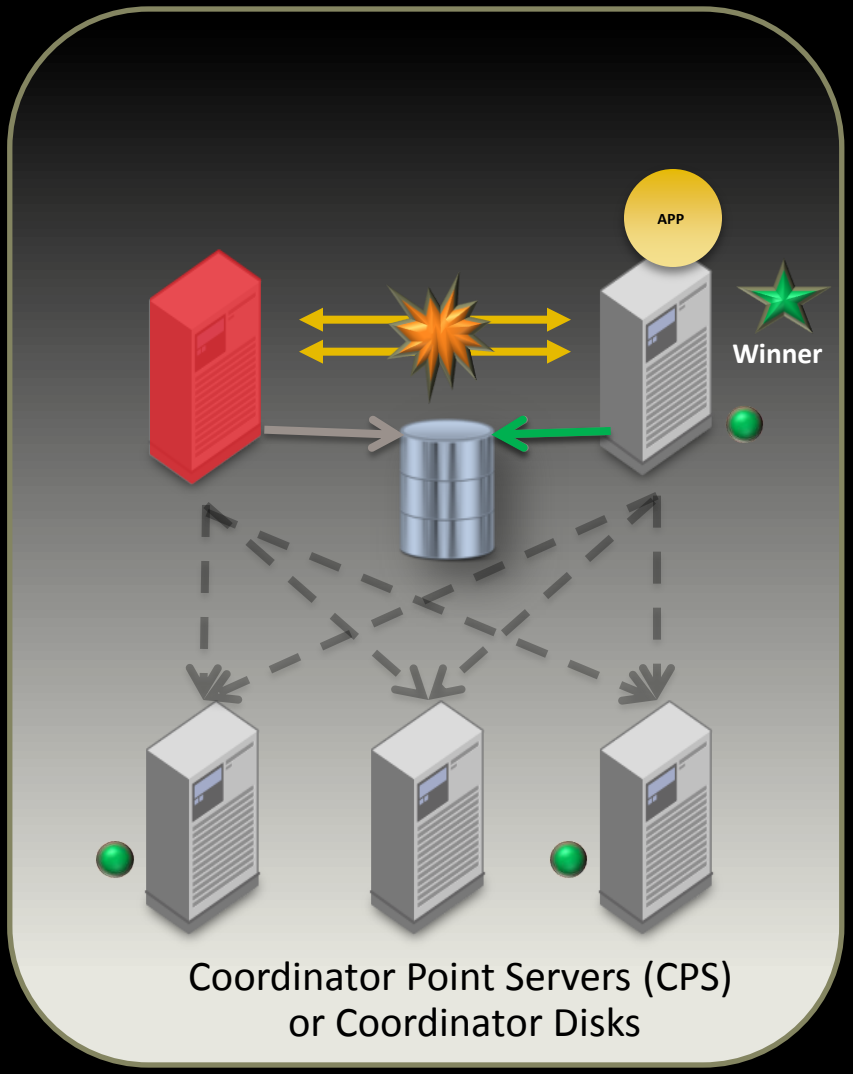

Fencing Priority

- Can be set on either the cluster node or the Service Group.
- When a fencing event occurs, the preferred fencing value is tabulated for the node to race.
- The priority is calculated by combining the priority for all nodes in the subcluster
- The lower the priority the longer the racer node will have to wait before attempting to contact the fencing coordination points.
- Priority fencing is gives priority but is not a guarantee a node will win.

Application Management

How We Do It

Centralized Management, Visibility, Reporting & Notification (VOM)

Application Management	Application Availability	Business Continuity
• Single node, either physical or virtual • App start, stop & restart • No local failover	• Local failover for apps • Customize based on local requirements • Cluster file system support	• Volume level and File level replication • Provides app disaster recovery • Ensure DR availability with Firedrill

Veritas Cluster Server

Features & Business Values

IT Challenge

How to effectively manage multi-tier services?

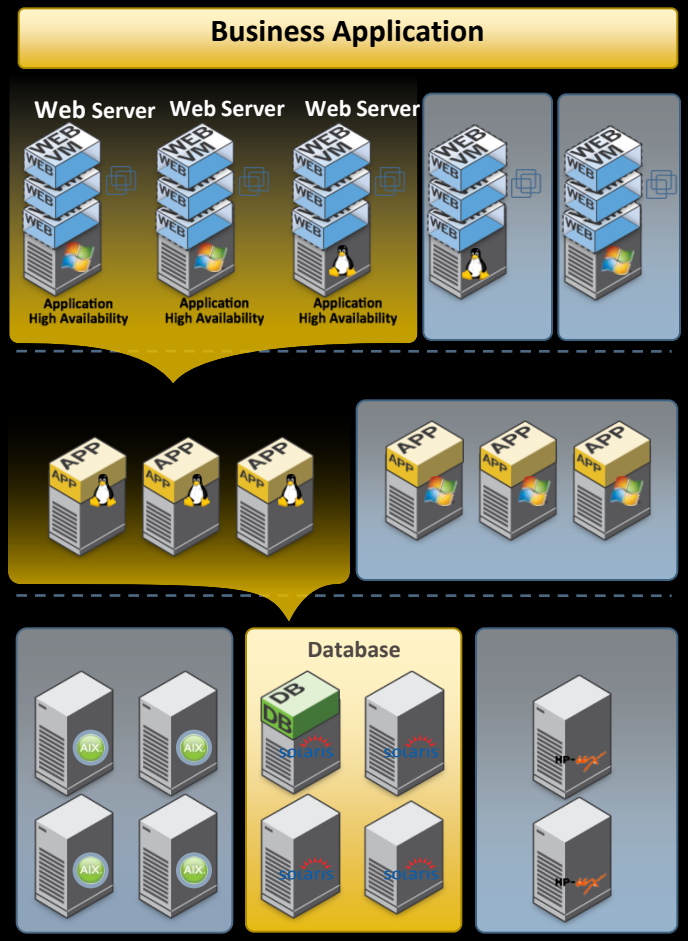

- Manage dependencies
- Service start/stop
- Service availability
- Disaster Recovery

Multi-Tier Applications & Services

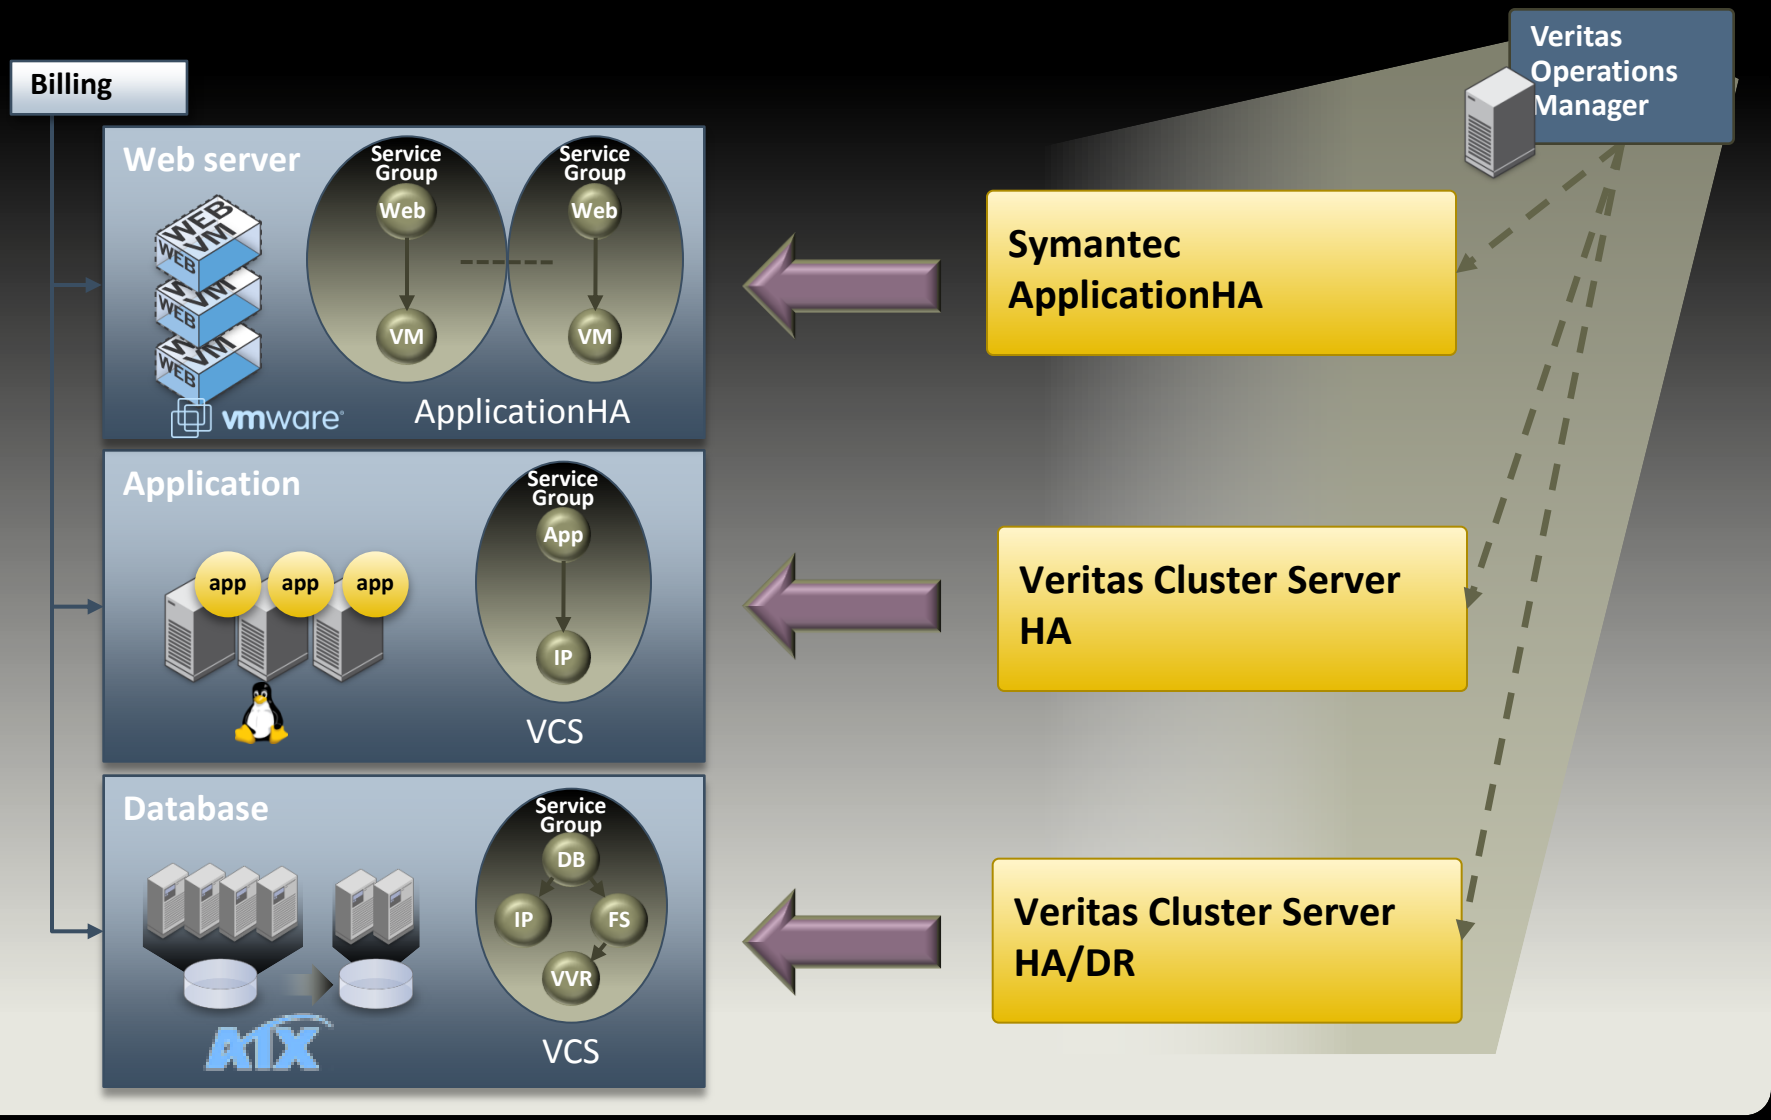

Transition to Virtual Business Service Seamlessly

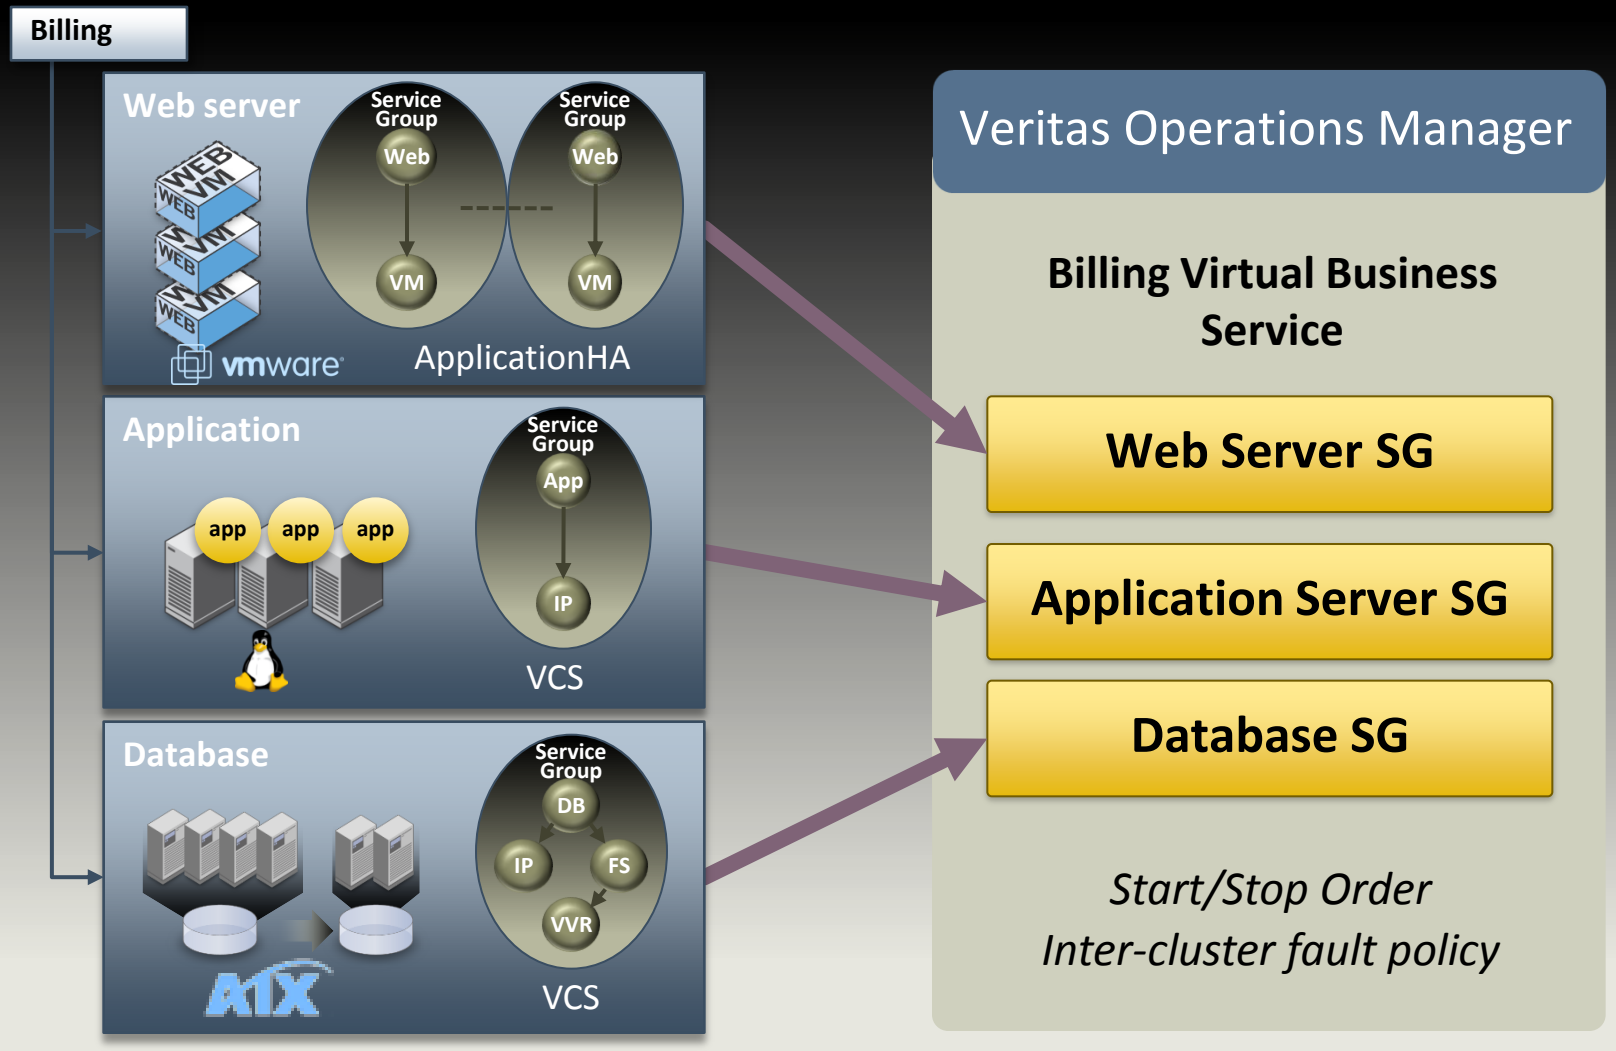

How Virtual Business Service Works

Start/Stop Orchestration

Communication Flow Between Clusters

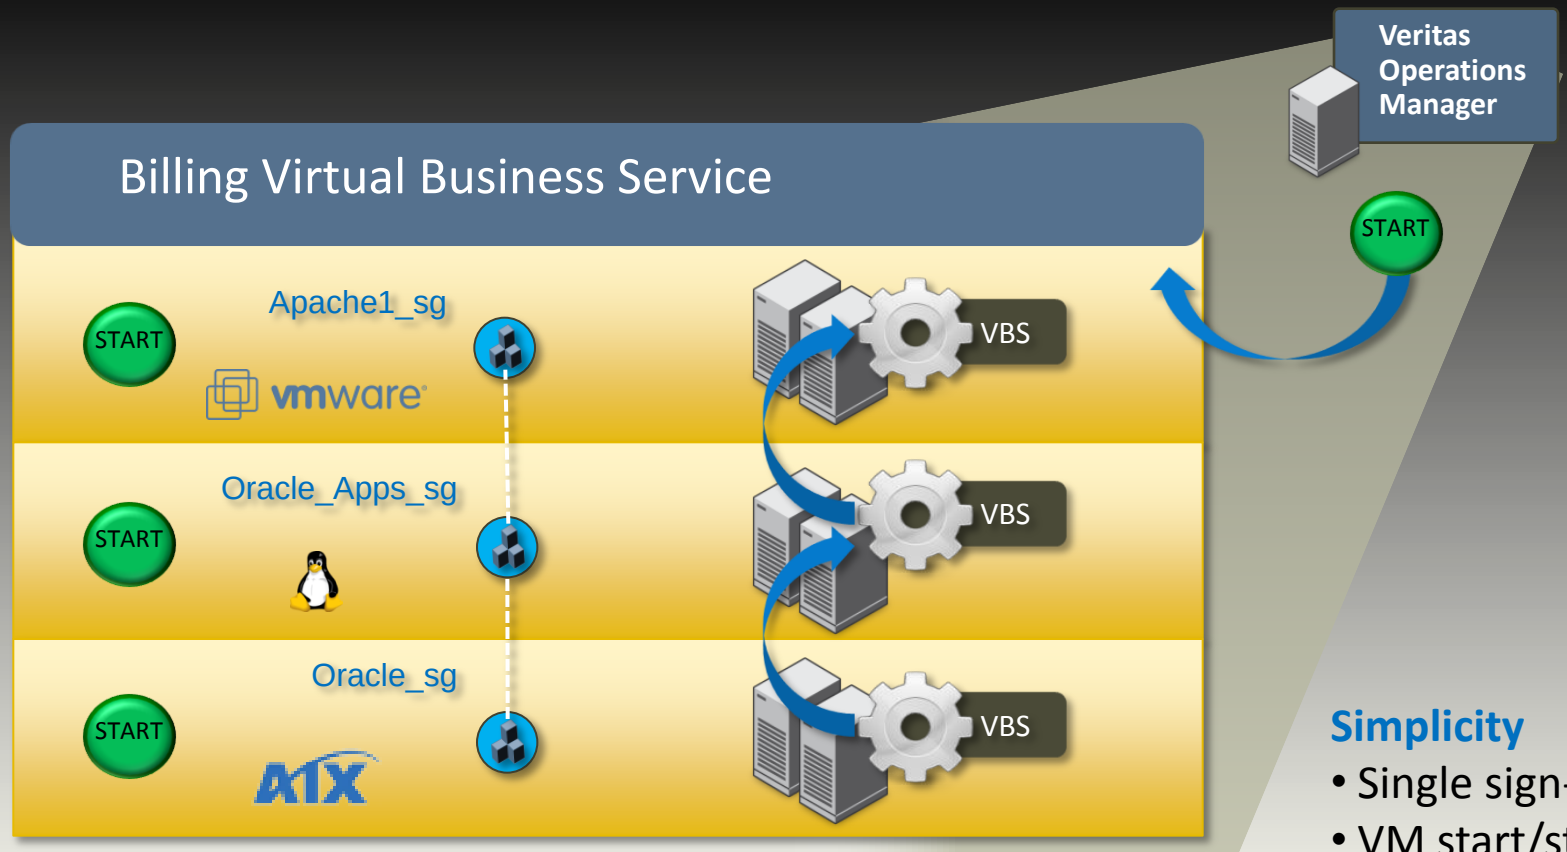

Simplicity

- Single sign-on
- VM start/stop
- Also through CLI

How Virtual Business Service Works

High Availability

Fault Propagation Between Clusters

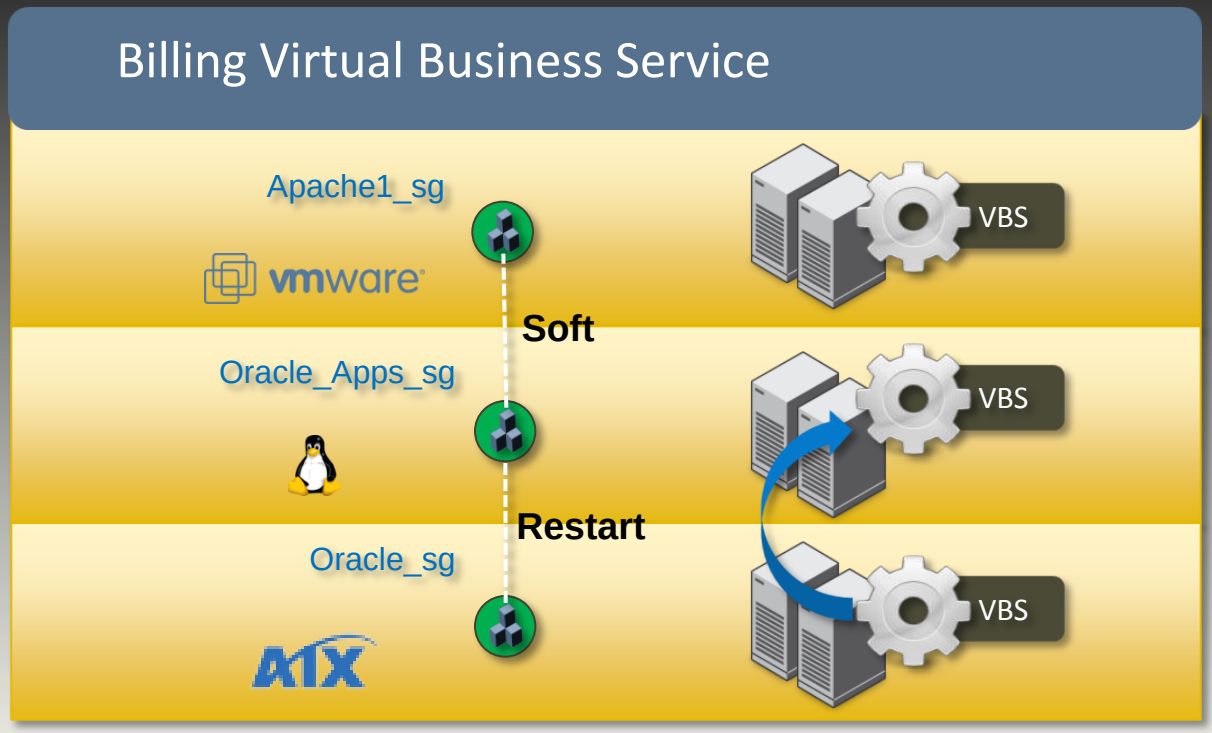

Configurable Fault Reaction

- Soft
- Firm
- Restart

Automatic Fault Propagation

- No external brain
- No heartbeats
- Secure

How Virtual Business Service Works

Disaster Recovery

Leveraging VCS Global Clusters for Service Level DR

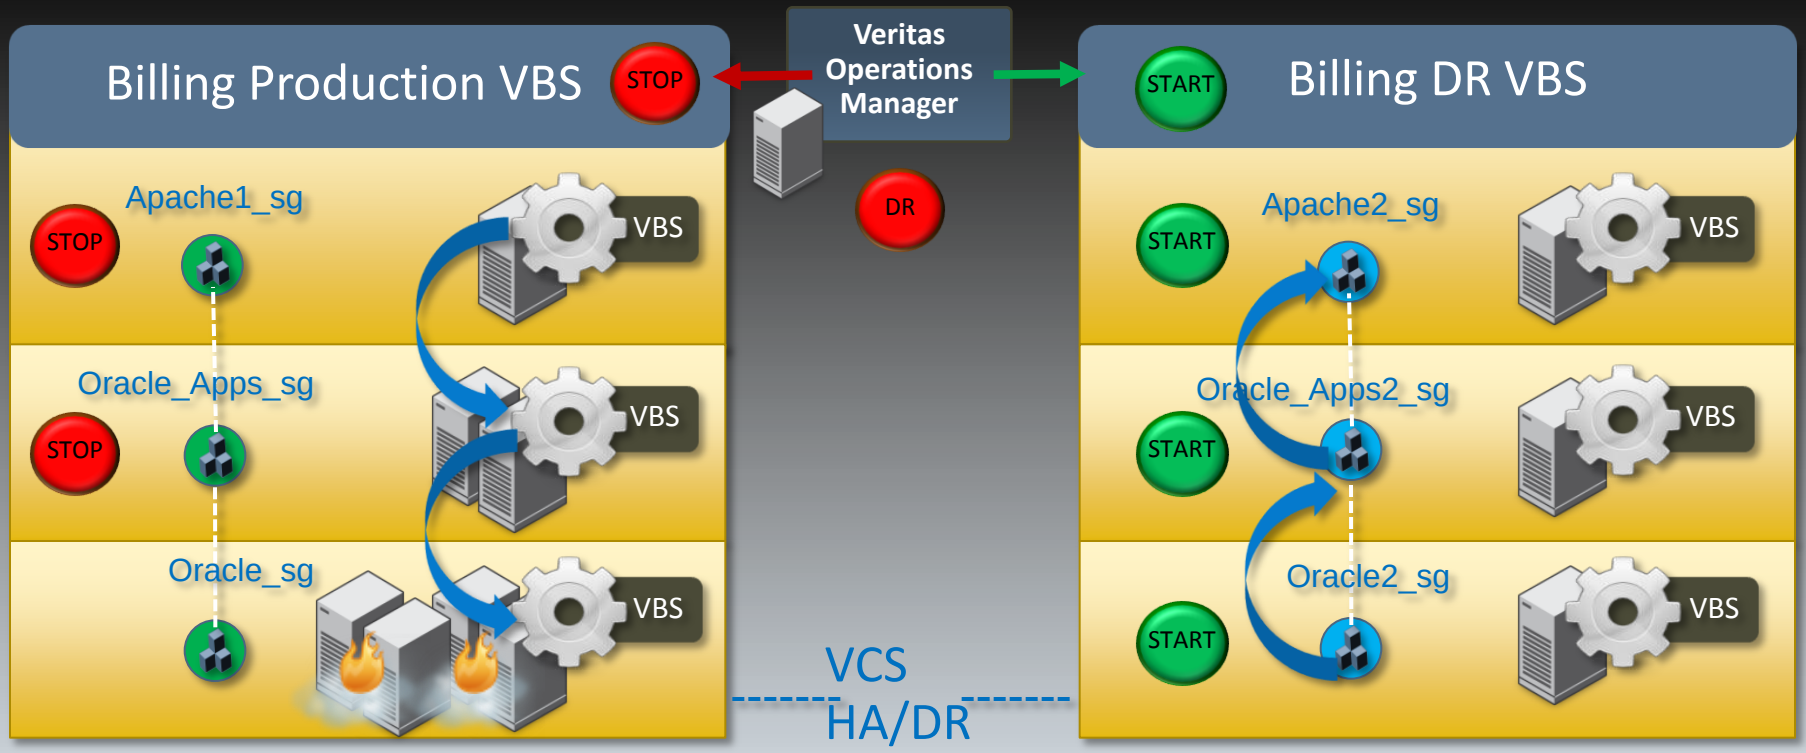

Flexibility

- Production & DR can have different configurations

No new infrastructure

- Leverages existing Global Cluster for DR of entire service

Virtual Business Service

Unparalleled Support Matrix

Operating Systems	Virtualization	High Availability
Solaris	VMware	VCS 5.1, 6.0
HP-UX	OracleVM, Zones	VCS 5.1 SP1 (Unix)
AIX	WPAR, LPAR	VCS 5.1 SP2 (Win)
Linux	KVM	ApplicationHA 5.1 SP2
Windows	Hyper-V	ApplicationHA 6.0

Databases	Applications	Replication
Oracle	Informatica	Veritas Replicator
Sybase	WebSphere	Oracle Dataguard
SAP	Microsoft Exchange	EMC SRDF
DB2	WebLogic	NetApp SnapMirror
MySQL	PeopleSoft	

Application Availability with Virtualization

Veritas Cluster Server & Symantec ApplicationHA

Application Availability in Virtualized Environments

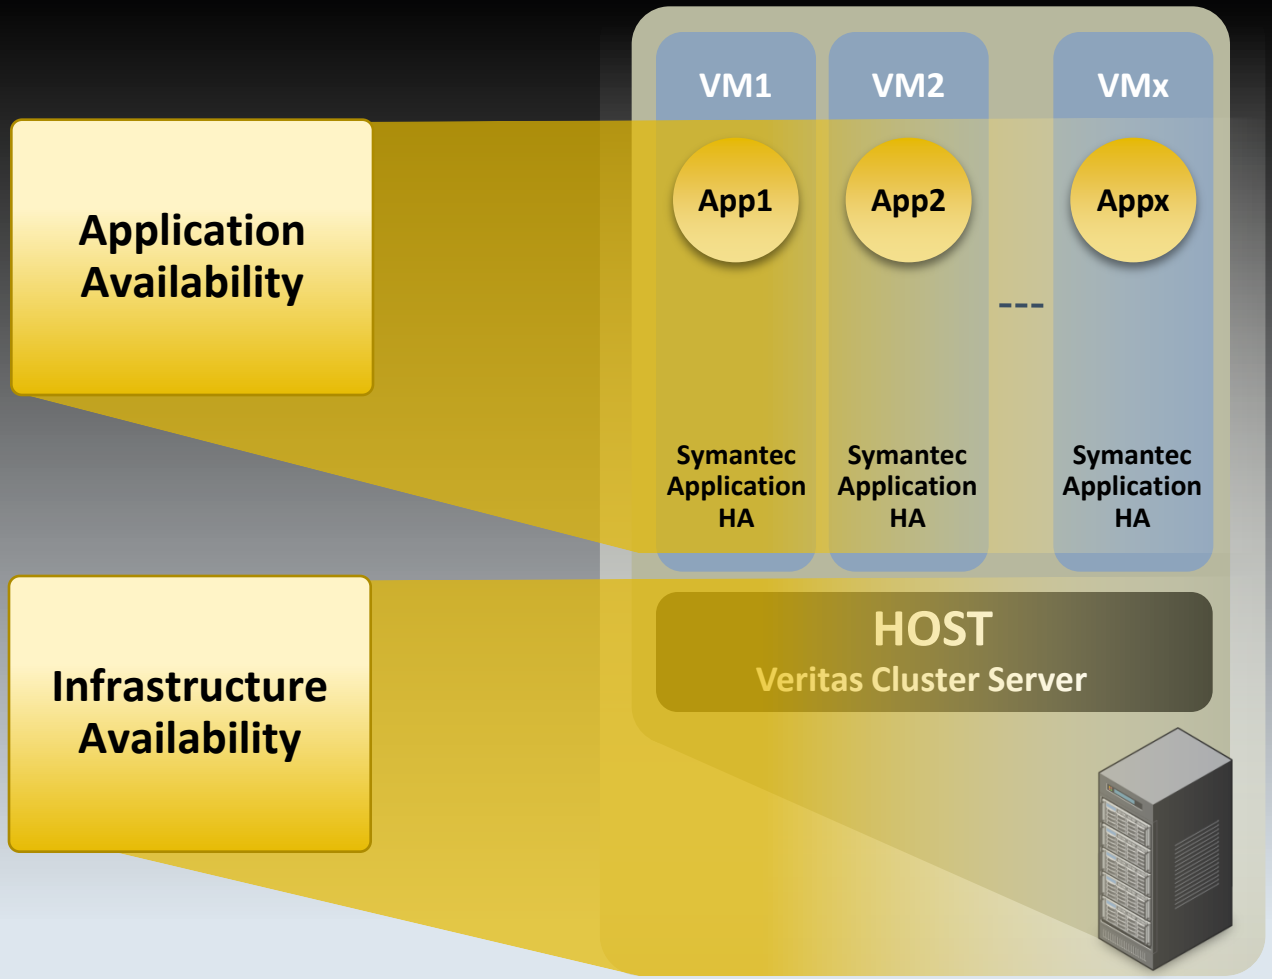

ApplicationHA Support

- Solaris LDOM
- Red Hat KVM
- AIX DLPAR

Symantec ApplicationHA

VMware Virtual Machines

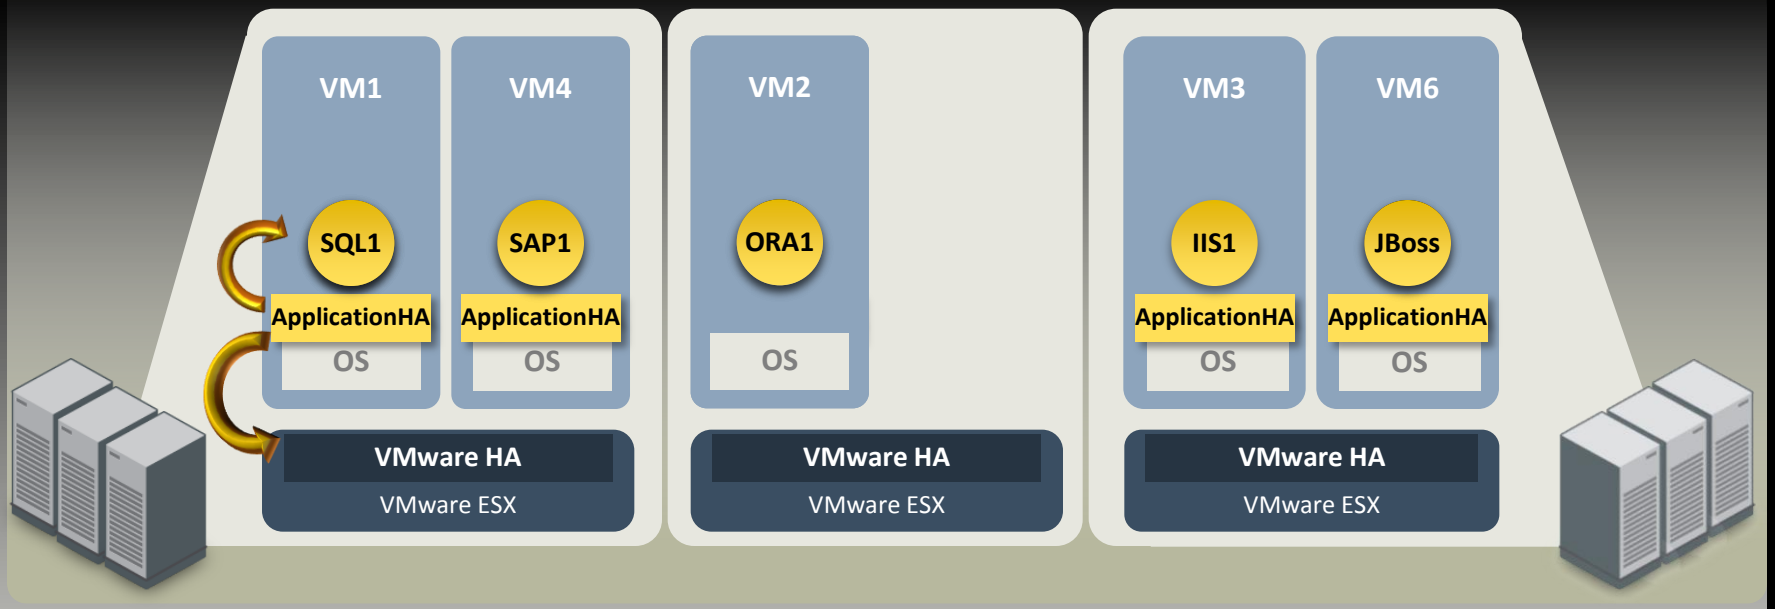

- Configured on Windows and Linux VMware Virtual Machines
- Enables Wizard Based Auto Configurations
- Provides Application protection within the Virtual Machine
- Managed through a vSphere plug-in
- Works in conjunction with VMwareHA
- Supports VMware vMotion

Capabilities

1. Without ApplicationHA VMwareHA does not have visibility into the VM to determine the state of the Application
2. ApplicationHA monitors the Application and can restart when there is an issue
3. If there continues to be an issue, Communicate with VMwareHA to Restart the Virtual Machine

ApplicationHA for application visibility and management

Simplify Application Management

- Manage app dependencies
- Start/Stop/Status of in-guest app

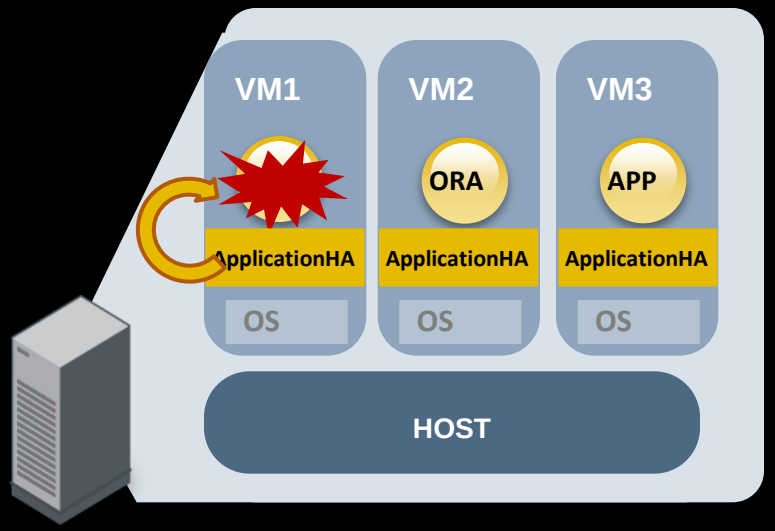

Fault detection and remediation

- Detailed application monitoring
- Application restart to remediate faults; no App/VM failover

Easy Configuration

- Automatic application discovery
- App wizards for configuring HA
- No private interconnects

Simplified Management

- In-guest app visibility through Veritas Operations Manager
- Fully integrated with Virtual Business Services for multi-tier application management

VCS + ApplicationHA for complete High Availability

Modular monitoring

- In guest App monitoring by ApplicationHA
- VM/Infrastructure monitoring by VCS
- Collaborative application aware VM monitoring

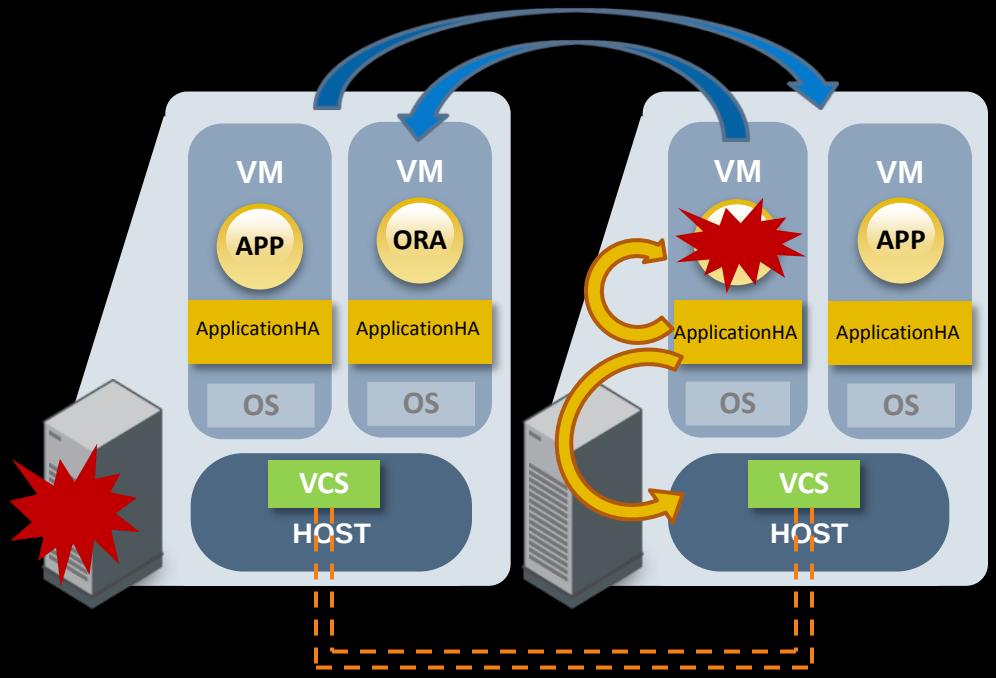

Coordinated application recovery

- Application restart
- VM restart on same server
- VM failover to standby server

Simplified Management

- End-to-end App/VM visibility through Veritas Operations Manager
- Clustering only at host level
- Internal communication between VCS and ApplicationHA for coordinated fault remediation

VCS Guest Cluster for 24x7 Application Availability

24x7 Intelligent Monitoring by VCS

- Detailed application awareness
- Infrastructure/VM Availability
- Traditional clustering

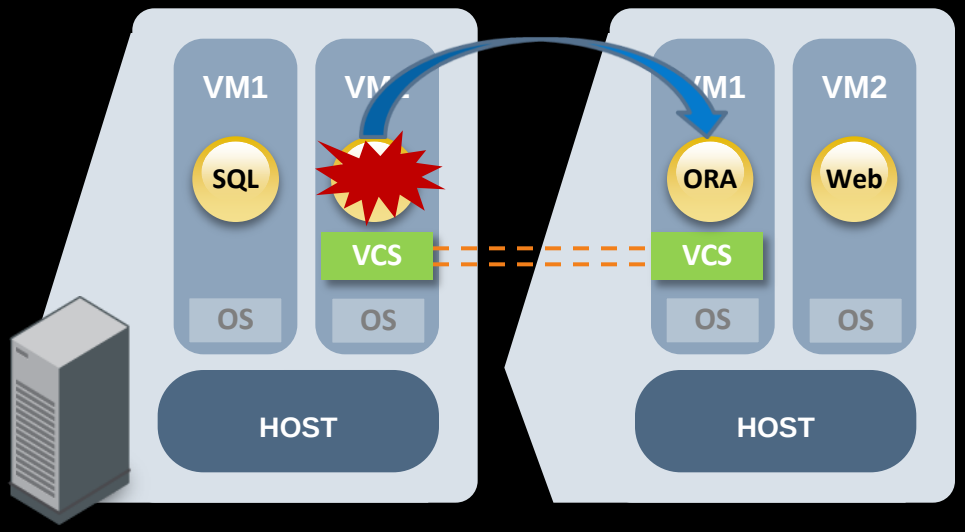

Rapid Application Recovery

- Fast failover to standby VM
- Integrated with SFCFS

Minimize Application Downtime

- Protection against OS corruption
- No downtime for OS patching

Simplified Management

- App visibility through Veritas Operations Manager
- Fully supported with Virtual Business Services

Thank you!

Anthony Herr

Anthony_Herr@symantec.com