



SQL Server High Availability: After Virtualization

SQL PASS Virtualization Virtual Chapter
September 11, 2013

About David Klee



vEXPERT 2013



Database Administrator 2008
Database Administrator on SQL Server® 2005
Database Developer 2008



- Founder and Chief Architect, Heraflux Technologies
- Teammate, Linchpin People
- Focus on SQL Server, cloud, infrastructure technologies
- Speaking at PASS Summit 2013

@kleegeek
 davidklee.net
 blog.heraflux.com
 gplus.to/kleegeek
 linkedin.com/a/davidaklee





About Heraflux Technologies

- Consulting services focused around business-critical data
- Areas of Focus:
 - Health and Stability
 - Performance Optimization
 - Virtualization and Cloud Enablement
 - Infrastructure Architecture
 - Risk Assessments
 - High Availability
 - Disaster Recovery
 - Capacity Management
 - Training and Education
- www.heraflux.com



© 2013 Heraflux Technologies



Agenda

- Goal: Less late night phone calls for you! How?
- Define High Availability (HA)
- Review SQL Server HA features
- Virtualization HA features
- How to select and combine the right features
- Virtualization-layer HA configuration tips



© 2013 Heraflux Technologies

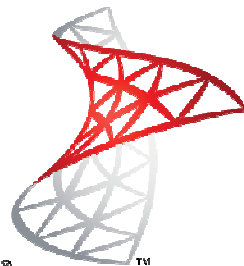
What is High Availability?

- "...system design approach and associated service implementation that ensures a prearranged level of operational performance that will be met during a contractual measurement period." – Wikipedia at <http://tinyurl.com/mlhud8>
- Distinguish between planned and unplanned downtime
- Metrics to care about:
 - Mean Time to Recovery (MTTR)
 - Recovery Time Objective (RTO)
- High Availability is not Disaster Recovery



SQL Server High Availability Options

- No HA (stand-alone instance)
- Traditional shared-storage clustering with WSFC
- Database-level mirroring
- AlwaysOn Availability Groups
- Replication
- Log shipping



Microsoft®
SQL Server®

Why Use SQL Server HA Features?

▪ Pros

- Reduce downtime for unplanned outages
- Reduce downtime for planned outages, such as maintenance or upgrades
- Reduce / eliminate single points of failure



▪ Cons

- Complexity
- Application compatibility
- Cost / Budget
- Internal supportability
- Potential performance impact



Virtualization High Availability

- VMware vSphere and Microsoft Hyper-V
- Use VM HA to buffer hardware component failures
- You probably already have it!
- Host-level HA features - resiliency
 - VMware vSphere High Availability (HA)
 - Microsoft Hyper-V failover cluster & SCVMM
- VM-level HA features
 - VMware vSphere Distributed Resource Scheduler (DRS)
 - Automatic placement of VMs on remaining hosts for resource maximization
 - VMware Fault Tolerance (FT) *

Features

✓ Turn On vSphere HA

vSphere HA detects failures and provides rapid recovery for the virtual machines running within a cluster. Core functionality includes host and virtual machine monitoring to minimize downtime when heartbeats cannot be detected.

vSphere HA must be turned on to use Fault Tolerance.

✓ Turn On vSphere DRS

vSphere DRS enables vCenter Server to manage hosts as an aggregate pool of resources. Cluster resources can be divided into smaller resource pools for users, groups, and virtual machines.

vSphere DRS also enables vCenter Server to manage the assignment of virtual machines to hosts automatically, suggesting placement when virtual machines are powered on, and migrating running virtual machines to balance load and enforce resource allocation policies.

vSphere DRS and VMware EVC should be enabled in the cluster in order to permit placing and migrating VMs with Fault Tolerance turned on, during load balancing.

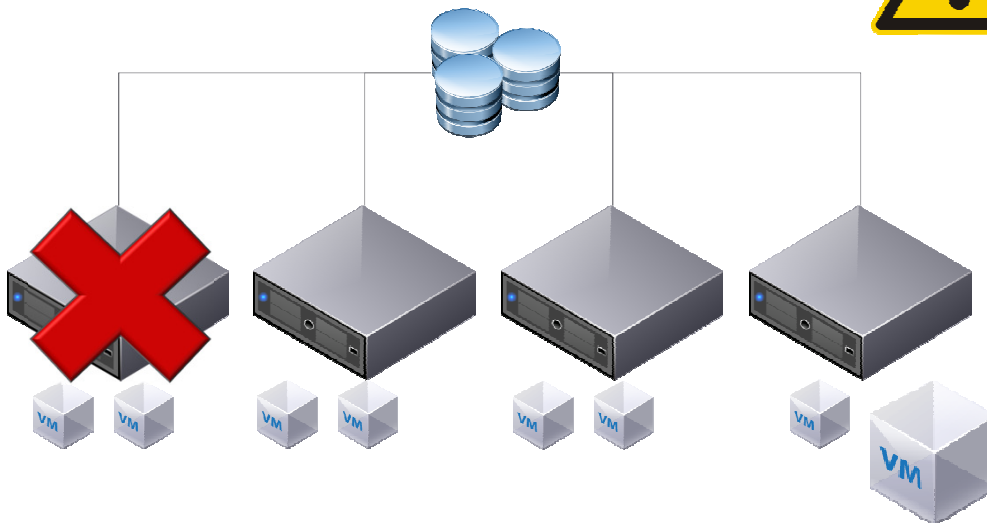
What does VM HA *not* do?

- Will not buffer against required reboots
- Will not automatically reboot VMs that have guest failures (i.e. BSOD)
 - But additional features or products can help you!
 - VMware VM HA Heartbeat even includes a screenshot of the console before it resets your VM
 - Third-party utilities can help monitor for VM OS failures
- Will not *eliminate* downtime from hardware failures
 - VMware Fault Tolerance – only VMs with 1vCPU
- Will not completely protect you from the architecture
 - i.e. blade chassis failure, path down, SAN failure

VM Monitoring Status
 VM Monitoring restarts individual VMs if their VMware tools heartbeats are not received within a set time. Application Monitoring restarts individual VMs if their VMware tools application heartbeats are not received within a set time.

VM Monitoring:

Virtualization High Availability



What Level of Protection?

- Criticality of the application?
- What is the cost of an outage?
- Planned downtime window duration?
- Unplanned downtime MTTR for all reasonable scenarios?
- Organizational downtime SLAs / RTOs?
- Available hardware in reserve for host failures?
- Edition of SQL Server / budget limitations?



Your PC ran into a problem and needs to restart. We're just collecting some error info, and then we'll restart for you. (0% complete)

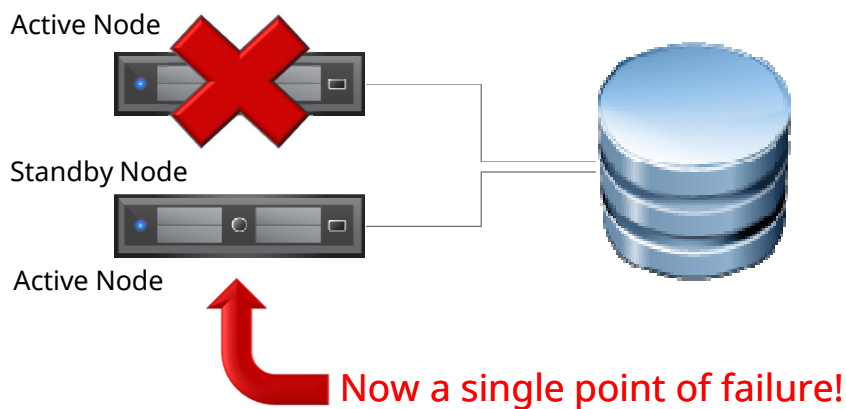
If you'd like to know more, you can search online later for this error: HAL_INITIALIZATION_FAILED

How To Select The Right Level

- Let the demand and numbers drive the technology!
- Examine the criticality of the system
- Determine unplanned outage MTTR / RTO
- Determine average window of time for planned outages
- Examine how you normally do things today
 - Is it time to overhaul your processes?
- Don't overcomplicate the design!

Reduce Risk

- Single best example from my years as a DBA...
Two-Node SQL Server Cluster

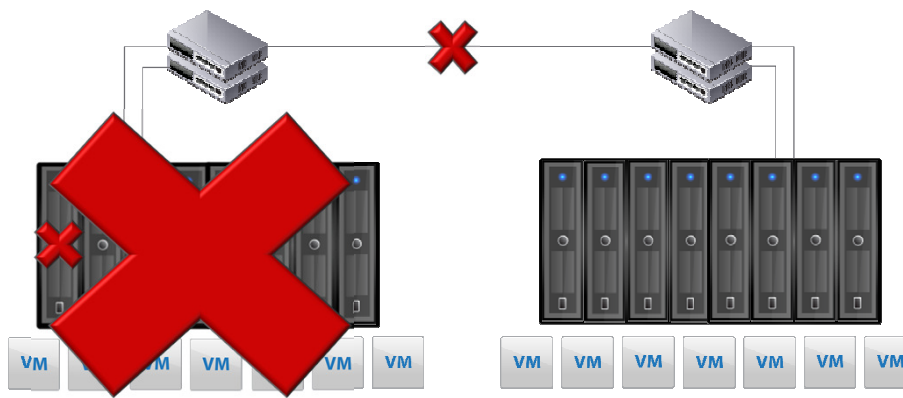


Common Scenarios

- VM HA + Single instance SQL Server for pre-production
- VM HA + SQL Server Cluster for production (pre-v2012)
- VM HA + SQL Server AlwaysOn Availability Groups (v2012+)
- VM HA + Single instance SQL Server for production Tier-2 applications (i.e. corporate antivirus, utility, etc. DBs)
- Tiers of criticality and tested HA solutions
 - Tier 1 (30 seconds): VM HA + Sync Availability Groups
 - Tier 2 (5 minutes): VM HA + Async Availability Groups
 - Tier 3 (30 minutes): VM HA Single Instance

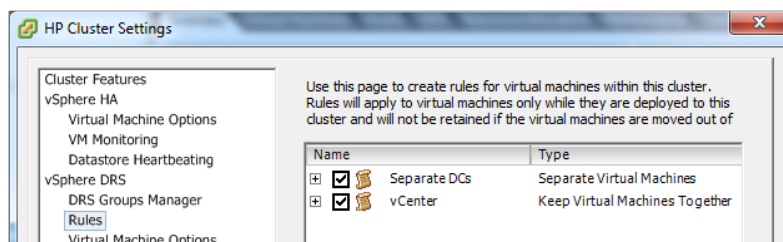
Failure Domains

- N+1 physical hosts does not always work
- Example: 16-node virtualized cluster – two blade chassis



Virtual Machine Placement Rules

- Pin VMs to certain hosts
 - Ex. vCenter, SCVMM, hardware dependency, SSD read-caching
- Keep VMs together
 - Ex. Application middle tier & database servers
- Separate VMs
 - Ex. Multiple nodes in a SQL Server Availability Group



Virtualization HA Configuration Tips

- Use cluster-level affinity / anti-affinity rules to group or separate critical VMs
- Set VM-level startup options
 - Ex. Database server before the application server
- Set startup priorities so business-critical VMs come up before secondary priority VMs
- Ensure that you have the right amount of spare resources to handle host failure scenarios

Demonstration

VMware vSphere HA Configuration

Conclusions

- Augment your SQL Server HA strategy with VM-level HA to reduce the hardware failure risk
- Enable VM-level HA through the hypervisor cluster-wide
- Ensure you have available resources to handle host failures
- Keep the architecture as simple as possible
- Work with your infrastructure admins to ensure no (reasonable) single points of failure
- ***Better resiliency means less unplanned late nights for you (and outages for your users)!***



