

云计算与虚拟化技术



第8讲: High Availability

阮晓龙

13938213680 / rxl@hactcm.edu.cn

<http://cloud.xg.hactcm.edu.cn>
<http://www.51xueweb.cn>

河南中医药大学信息管理与信息系统教研室
信息技术学院网络与信息系统科研工作室

2019.4

讨论提纲

- Introducing High Availability
 - Understanding the Layers of High Availability
 - Clustering VMs
- Implementing vSphere High Availability
 - Understanding vSphere High Availability Clusters
 - Understanding vSphere High Availability's Core Components
 - Enabling vSphere HA
 - Configuring vSphere High Availability
 - Managing vSphere High Availability
- Planning for Business Continuity



1.Introducing

1.1Understanding the Layers of High Availability

- Even in non-virtualized environments, there are multiple ways to achieve high availability for Operating System (OS) instances and applications. When you introduce virtualization into the mix with vSphere, you gain additional methods of providing high availability.
 - At each layer are tools and techniques for providing high availability and business continuity:
 - At the Application layer, options include Oracle Real Application Clusters (RAC), or simply application resiliency through a micro-services architecture.
 - At the OS layer, solutions include OS clustering functionality, such as Windows Server Failover Clustering (WSFC).



1.Introducing

1.1 Understanding the Layers of High Availability

- Even in non-virtualized environments, there are multiple ways to achieve high availability for Operating System (OS) instances and applications. When you introduce virtualization into the mix with vSphere, you gain additional methods of providing high availability.
 - At each layer are tools and techniques for providing high availability and business continuity:
 - The Virtualization layer offers a number of features for high availability, including vSphere High Availability (HA), vSphere Fault Tolerance (FT), and stretched (metro) storage clusters.
 - At the Physical layer, high availability is achieved through redundant hardware—multiple network interface cards (NICs) or host bus adapters (HBAs), multiple storage area network (SAN) switches and fabrics, multiple paths to storage, multiple controllers in storage arrays, redundant power supplies, and so forth.




1.Introducing

1.1 Understanding the Layers of High Availability

FIGURE 7.1

Each layer has its own forms of high availability.



1.Introducing

1.1Understanding the Layers of High Availability

- Each of these technologies or techniques has its own strengths and weaknesses. For example:
 - providing redundancy at the Physical layer is great, but it doesn't help with failures at the Application layer.
 - Conversely, protecting against application failures won't help much if the underlying hardware isn't redundant.
- As you set forth to establish high availability for your virtualized workloads, keep in mind that there is no "one size fits all" solution.
- Use the right tool for the job based on your specific requirements.



1.Introducing

1.1 Understanding the Layers of High Availability

- On vSphere, four key technologies or techniques that help provide high availability:
 - OS clustering in Microsoft Windows
 - ESXi host clustering using vSphere HA
 - Virtual machine (VM) mirroring using vSphere FT
 - vSAN stretched storage clustering



1.Introducing

1.2Clustering VMs

- Because Windows Server is widely used in corporate and enterprise datacenters today, it's quite likely that you've been asked to create or support a Windows-based cluster.
- There are two primary ways to use clustering to provide high availability for Windows Server:
 - Network Load Balancing (NLB) clustering
 - Windows Server Failover Clustering (WSFC)




1.Introducing

1.2Clustering VMs

FIGURE 7.2

An NLB cluster can contain up to 32 active nodes (only 5 are shown here), and traffic is distributed equally across each available node. The NLB software allows the nodes to share a common name and IP address that is referenced by clients.




1.Introducing

1.2Clustering VMs

FIGURE 7.3

Server clusters are best suited for applications and services like SQL Server, DHCP, and so on, which use a common dataset.




2.Implementing vSphere HA

2.1 Understanding vSphere HA Clusters

- The vSphere High Availability feature, or more simply, vSphere HA, is designed to provide an automatic restart of the VMs that were running on an ESXi host at the time it became unavailable.

FIGURE 7.10

vSphere Availability provides an automatic restart of VMs that were running on an ESXi host when it failed.



2.Implementing vSphere HA

2.1Understanding vSphere HA Clusters

- ❑ vSphere HA primarily targets ESXi host failures, but it can also protect against VM failures and respond to storage connectivity issues. In all cases, vSphere HA uses a restart of the VM as the mechanism for addressing the detected failure. This means there is a period of downtime when a failure occurs. Unfortunately, you can't calculate the exact duration of the downtime because it is unknown ahead of time how long it will take to boot a VM or a series of VMs.
- ❑ vSphere HA might not provide the same level of high availability found in other high-availability solutions. when a failover occurs between ESXi hosts as a result of the vSphere HA feature, there is a slight potential for data loss and/or filesystem corruption because the VM was immediately powered off when the server failed and then brought back up minutes later on another server.



2.Implementing vSphere HA

2.2vSphere HA's Core Components

- HA uses a VMware-developed tool known as Fault Domain Manager (FDM).
- FDM offers a few significant improvements:
 - FDM uses a master/slave architecture that does not rely on primary/secondary host designations.
 - FDM uses both the management network and storage devices for communication.
 - FDM supports IPv6.
 - FDM addresses the issues of both network partition and network isolation.



2.Implementing vSphere HA

2.2vSphere HA's Core Components

- When vSphere HA is enabled, the vSphere HA agents participate in an election to pick a vSphere HA master. The vSphere HA master is responsible for the following key tasks within a vSphere HA-enabled cluster:
 - Monitors slave hosts and will restart VMs in the event of a slave host failure.
 - Monitors the power state of all protected VMs. If a protected VM fails, the vSphere HA master will restart the VM.
 - Manages the list of hosts that are members of the cluster and manages the process of adding and removing hosts from the cluster.



2.Implementing vSphere HA

2.2vSphere HA's Core Components

- When vSphere HA is enabled, the vSphere HA agents participate in an election to pick a vSphere HA master. The vSphere HA master is responsible for the following key tasks within a vSphere HA-enabled cluster:
 - Manages the list of protected VMs. It updates this list after each user-initiated power-on or power-off operation. These updates are at the request of vCenter Server, which requests the master to protect or unprotect VMs.
 - Caches the cluster configuration. The master notifies and informs slave hosts of changes in the cluster configuration.



2.Implementing vSphere HA

2.2vSphere HA's Core Components

- When vSphere HA is enabled, the vSphere HA agents participate in an election to pick a vSphere HA master. The vSphere HA master is responsible for the following key tasks within a vSphere HA-enabled cluster:
 - The vSphere HA master host sends heartbeat messages to the slave hosts so that the slave hosts know the master is alive.
 - Reports state information to vCenter Server. vCenter Server typically communicates only with the master.



2.Implementing vSphere HA

2.2vSphere HA's Core Components

- As you can see, the role of the vSphere HA master is quite important.
 - For this reason, if the existing master fails a new vSphere HA master is automatically elected.
 - The new master will then take over the responsibilities listed here, including communication with vCenter Server.



2.Implementing vSphere HA

2.2vSphere HA's Core Components

- Once an ESXi host in a vSphere HA-enabled cluster elects a vSphere HA master, all other hosts become slaves connected to that master. The slave hosts have the following responsibilities:
 - A slave host watches the runtime state of the VMs running locally on that host. Significant changes in the runtime state of these VMs are forwarded to the vSphere HA master.
 - vSphere HA slaves monitor the health of the master. If the master fails, slaves will participate in a new master election.
 - vSphere HA slave hosts implement vSphere HA features that don't require central coordination by the master. This includes VM health monitoring.



2.Implementing vSphere HA

2.3Enabling vSphere HA

- To implement vSphere HA, you must ensure all of these requirements are met:
 - All hosts in a vSphere HA-enabled cluster must have access to the same shared storage locations used by all VMs on the cluster. This could be any Fibre Channel, FCoE, iSCSI, NFS, or vSAN datastores used by VMs.
 - All hosts in a vSphere HA cluster should have an identical virtual networking configuration. If a new switch is added to one host, the same new switch should be added to all hosts in the cluster. If you are using a vSphere Distributed Switch (vDS), all hosts should be participating in the same vDS.



2.Implementing vSphere HA

2.3Enabling vSphere HA



- Enabling vSphere HA On VCSA
 - Create Datacenter
 - Create Cluster
 - Add ESXi Host
 - Enabling vSphere HA



2.Implementing vSphere HA

2.4Configuring vSphere HA

- After vSphere HA is enabled, there are a number of ways to configure the different components of vSphere HA. Grouped within the UI, these configuration items revolve around several key areas:
 - Failures and Responses
 - Proactive HA Failure and Responses
 - Admission Control
 - Heartbeat Datastore
 - Advanced Options



2.Implementing vSphere HA

2.4Configuring vSphere HA



- Configuring vSphere HA On VCSA
 - Detailing Failure and Response Settings
 - Getting Started with Proactive HA Failure and Responses
 - Configuring vSphere HA Admission Control
 - Heartbeat Datastores
 - Configuring vSphere HA Groups, Rules, Overrides, and Orchestrated VM Restart



2.Implementing vSphere HA

2.5Managing vSphere HA

- ❑ Many vSphere HA functions are calculated automatically, such as slot size, total number of slots, selection of hosts for datastore heartbeating, and the selection of the master/slave roles by FDM. Without proper exposure of these values, it would be difficult to properly manage vSphere HA and its operation.
- ❑ VMware included information about vSphere HA in the vSphere web client to make it easier to manage vSphere HA.



2.Implementing vSphere HA

2.5Managing vSphere HA



- Managing vSphere HA On VCSA
 - The Summary Area
 - Heartbeat Area
 - Configuration Issues Area
 - Datastores Under APD Or PDL



3.Planning for Business Continuity

- ❑ High availability is only part of the solution; it's one component in the bigger picture of business continuity. Business continuity is about ensuring that the business can continue operating in the face of a significant event.
- ❑ High availability deals with business continuity from a fairly narrow perspective: ensuring that the business can continue operating in the event of a physical server failure, an OS or application failure, or a network component failure. There are many types of failures that you must account for and protect against.



3.Planning for Business Continuity

- Discuss the following two primary types here:
 - Protecting against the loss of data due to equipment failure, software malfunction, or simple user error (such as deleting something by mistake, which most of us have done at one time or another).
 - Planning for disaster recovery in the event your entire datacenter is rendered unusable or unavailable.




3.Planning for Business Continuity

3.1Providing Data Protection

FIGURE 7.38

Running backup agents inside the guest OS can provide application- and OS-level integration, but not without some drawbacks.



3.Planning for Business Continuity

3.2Recovering from Disasters

- High availability makes up only half of the ability to keep your application/systems up in day-to-day operation. The **other half is disaster recovery**, which is the ability to recover from a catastrophic failure.
 - The risks posed by hurricanes, earthquakes, and other natural and manmade disasters underscore how important it is to establish a thoughtfully designed plan that you can execute with certainty.
 - Entire datacenters can be destroyed by one of these events, and even the datacenters that survive and keep functioning do not stay operational for long when generators run out of gas.



3.Planning for Business Continuity

3.3Using vSphere Replication

FIGURE 7.39

vSphere Replication can work between datacenters, as long as there is a network joining them.

