

TAC 9710 - Virtualizing a Windows Active Directory Domain Infrastructure

Chris Skinner

Technical Instructor
Education Services

VMware, Inc.



VMWORLD 2006

Why Virtualize Active Directory?

- Hardware Consolidation
- Test and Development
- Security Control

Hardware Consolidation

- Combine multiple, single use boxes
- Standardization – eliminate imaging problems
- Reduce Product Activation issues

Test and Development

- Policy Testing
- Schema Changes
- Migration/Upgrade testing
- Domain reconfigurations
- Deployment testing
- Disaster recovery planning

Security control

- Physical access
- Administrative delegations
- Separate applications from Active Directory databases

Supported Operating Systems

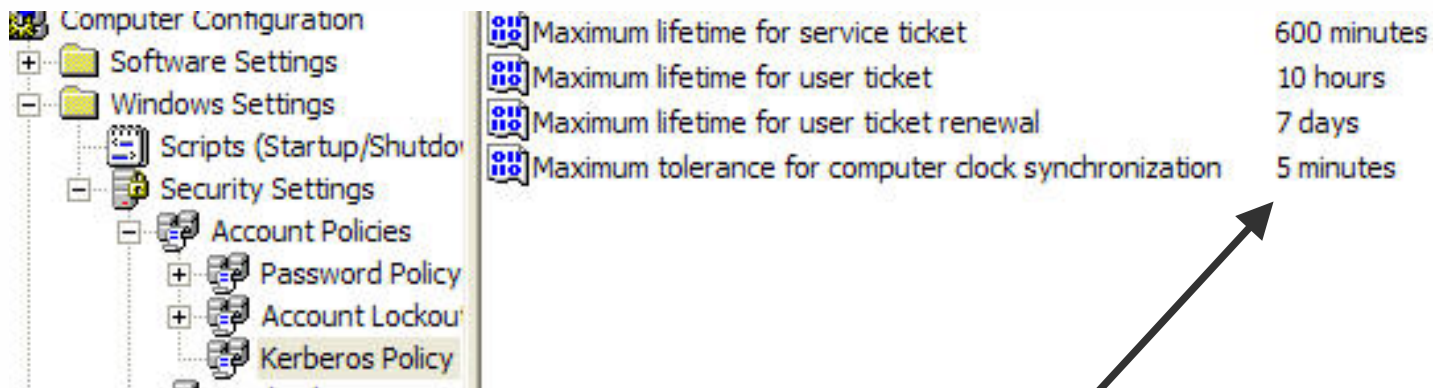
- Windows 2000
- Windows 2003 & R2
- Native Mode or Mixed Mode

Challenges to Virtualizing Active Directory

- Clock synchronization
- Network performance
- Multi-master replication model
- Security
- Potential single point of failure
- Disaster recovery

Clock Synchronization – Why So Important?

- Active Directory operations are critically time dependent
- MS Kerberos implementation allows a 5 minute tolerance
- File Replication Services (FRS) synchronizes scripts, database changes/updates, policies based, in part, on time-stamping



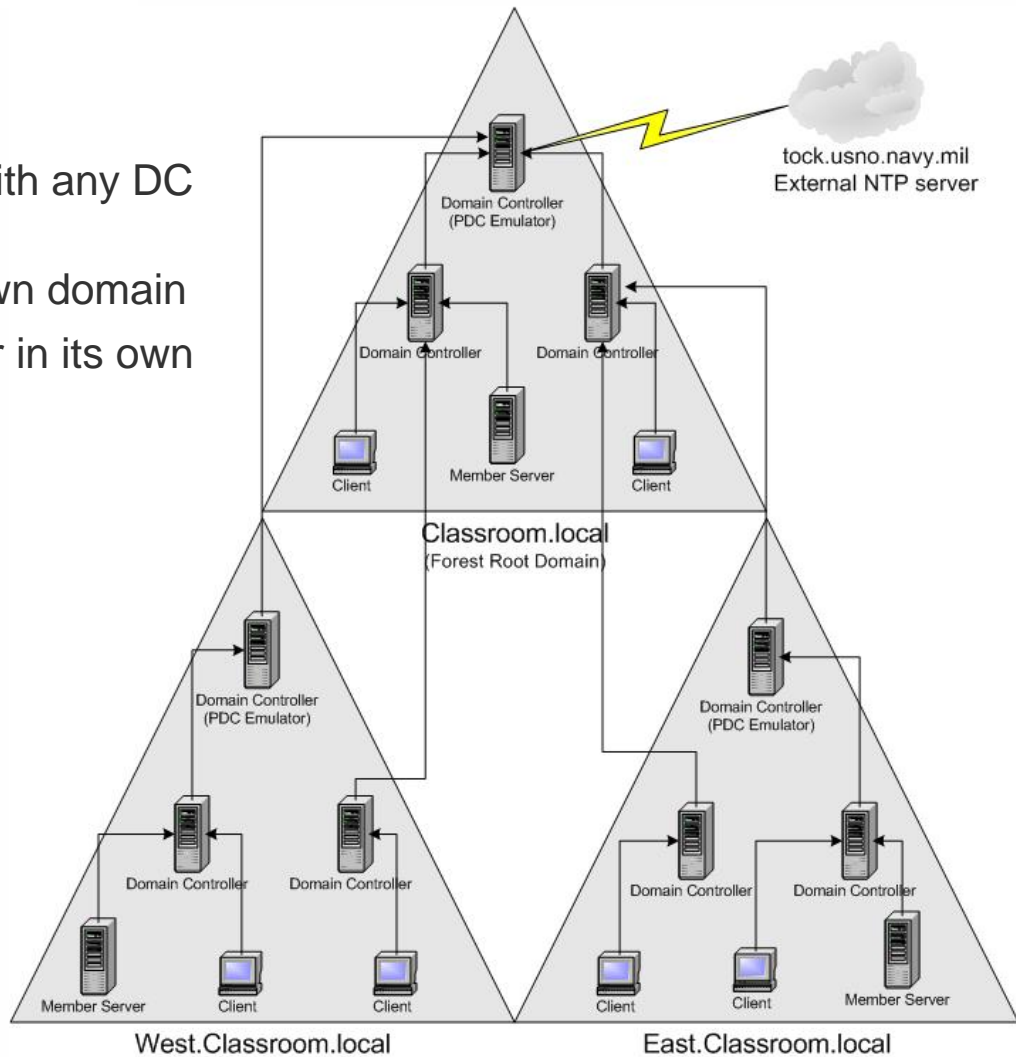
The image shows a screenshot of the Windows Group Policy console. On the left, a tree view is expanded to 'Computer Configuration > Windows Settings > Security Settings > Account Policies > Kerberos Policy'. On the right, four policies are listed with their corresponding values:

Maximum lifetime for service ticket	600 minutes
Maximum lifetime for user ticket	10 hours
Maximum lifetime for user ticket renewal	7 days
Maximum tolerance for computer clock synchronization	5 minutes

An arrow points from the bottom right towards the 'Maximum tolerance for computer clock synchronization' policy.

Time Server Hierarchies

- Child PDC emulators can sync with any DC in the parent domain
- Clients sync with any DC in its own domain
- DCs can sync with PDC emulator in its own domain or any DC in parent



Source: Microsoft Corporation

Clock Synchronization – Virtualization Issues

- No CPU cycles needed – none given!
- Clock drifts can be significant in a relatively short period
- Idle cycles in a virtual machine is an Active Directory domain's worst enemy
- How do you combat time synchronization issues?

More than a 28
minute drift!



Congratulations ! System time corrected by 1700 second(s)
is now operating as a time server on your network.
To customize for use with GPS, NTP4 or other
Reference Time Sources, please use the Settings tab.
Refer to the Log tab for synchronization messages.

OK

**During 18
hour test**

Clock Syncing – Option A – Using W32Time

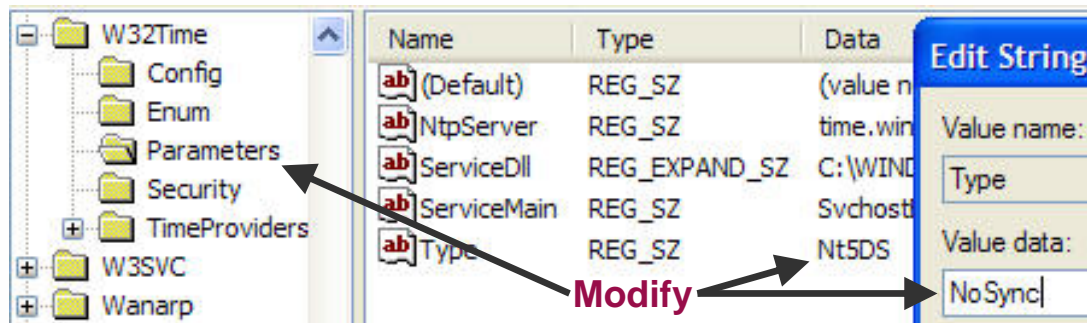
- Use Windows Time Service – NOT VMware Tools
 - Define an alternative external time source for “master” time server
1. Modify Registry settings on the PDC emulator for the forest root domain:
`HKLM\System\CurrentControlSet\Services\W32Time\Parameters`
 - Change **Type RED_SZ** value from **NT5DS** to **NTP**
 - Change **NtpServer** value from **time.windows.com,0x1** to an external stratum 1 time source, i.e. **tock.usno.navy.mil,0x1**
`HKLM\System\CurrentControlSet\Services\W32Time\Config`
 - Change **AnnounceFlags REG_DWORD** from **10** to **5**
 2. Stop and restart Time Service – **net stop w32time → net start w32time**
 3. Manually force update → **w32tm /resync /rediscover**

Name	Type	Data
(Default)	REG_SZ	(value not set)
AnnounceFlags	REG_DWORD	0x00000005 (5)
EventLogFlags	REG_DWORD	0x00000002 (2)

Name	Type	Data
(Default)	REG_SZ	(value not set)
NtpServer	REG_SZ	tock.usno.navy.mil,0x1
ServiceDll	REG_EXPAND_SZ	C:\WINDOWS\system32\w32time.dll
ServiceMain	REG_SZ	SvchostEntry_W32Time
Type	REG_SZ	NTP

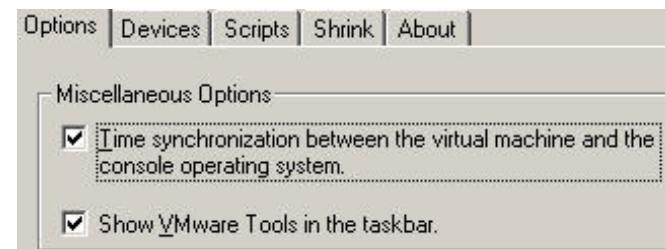
Clock Synching – Option B – Using VMware Tools

- Modify Windows Time Service – Use VMware Tools
 - Implement Domain Controllers Group Policy to modify registry:



- Enable ESX server NTP daemon to sync with external stratum 1 NTP source
 - VMware Knowledge Base ID# 1339
- Use VMware Tools Time Synchronization within the virtual machine

NOTE: VMware Tools time sync is designed to play “catch-up”, not slow down!

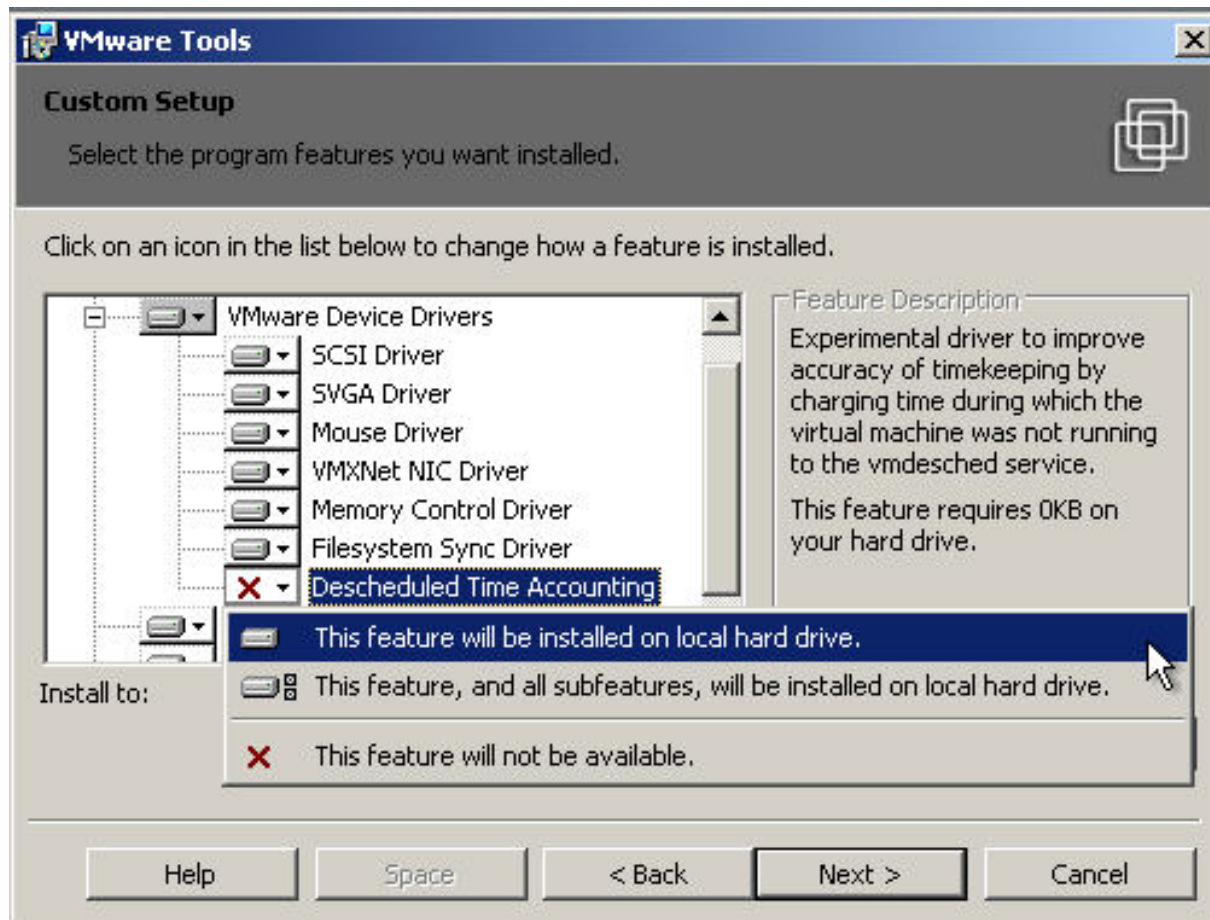


New Clock Syncing – Descheduled Time Accounting

- VMware Tools experimental component
- Custom component for ESX 3.x VMs
- Currently for uniprocessor Windows and Linux VMs
- Improved accuracy for guest OSes CPU time accounting
- Allows quicker “catch-up” of time for guest OS
- Launches a VMDesched thread or process

Clock Synching – Descheduled Time Accounting (2)

- Perform a Custom installation of VMware Tools in Windows guest OS



ESX 3.x/VC 2.x Security - VM Access Control

Hosts & Clusters

- Corporate
 - Domain Controllers
 - Headquarters
 - 192.168.32.53
 - 192.168.32.54
 - 192.168.32.57
 - East DCs
 - ChildDC3
 - ChildDC4
 - Root DCs
 - RootDC2
 - RootDC
 - West DCs
 - ChildDC1
 - ChildDC2
 - Win2kMS1

Domain Controllers

Virtual Machines | Hosts | Tasks & Events | Alarms | **Permissions** | Maps

User/Group	Role	Defined in
CLASSROOM\DC VM Admins	Virtual Machine Power User	This object
Administrators	Administrator	Hosts & Clusters

Select Users

Select users and groups to include in this role. You can also manually enter names and use the Check Names Feature to validate your entries against the directory.

Domain: CLASSROOM

Users and Groups

Show Users First [] Search []

Name
Administrator
Guest
krbtgt
SUPPORT_388945a0
DC VM Admins
DnsUpdateProxy

Add []

Users: []

Groups: CLASSROOM\DC VM Admins

Note: Separate multiple names with semicolons.

Check Names []

Ok [] Cancel []

Add their names to the Users and Groups list below. Assign them a role from the drop-down menu to the right.

Assigned Role

Selected users and groups can interact with the current object according to the role and privileges chosen below.

Virtual Machine Power User

- Global
- Folder
- Datacenter
- Datastore
- Network
- Host
- Virtual Machine
- Resource
- Alarms
- Scheduled Task
- Sessions
- Performance
- Permissions

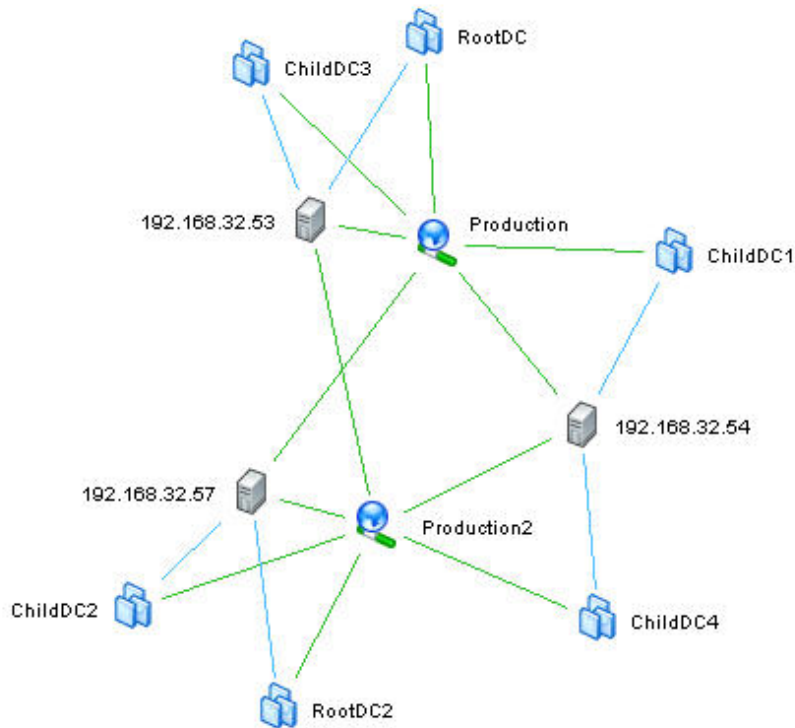
Description: All Privileges

Propagate to Child Objects

Transitioning from Physical to Virtual

- Start with a fresh system state backup for recovery
- Consider creating a dedicated virtual switch or virtual machine port group to isolate replication traffic
- Generally single processor virtual machines are adequate for domain controllers
- Create a separate virtual disk for Active Directory database, log files, and SYSVOL
- Validate inbound/outbound connections between physical and virtual machines
- Allow 24-48 hours for replication to complete
- Change the weight and/or priority of the DNS SRV records for virtual machines
- Monitor the logon requests to ensure virtual machines are successfully responding
- Decommission physical domain controllers

Network Connections



Create separate VM port groups connected to individual NICs

[Refresh](#)



Use the Maps view to verify network infrastructure

Map Relationships:

Custom Map

Host Options

- Host To VM
- Host To Network
- Host To Data
- Show All

VM to Network

VM to Data

Show All

Apply

Virtual Switch: vSwitch1

Virtual Machine Port Group

- Production
- 2 virtual machines | VLAN ID *
- Win2kM51
- ChildDC1

Physical Adapters

- vmnic1 1000 Full

Virtual Switch: vSwitch3

VMkernel Port

- VMotion
- 10.1.32.154

Virtual Machine Port Group

- Production2
- 1 virtual machines | VLAN ID *
- ChildDC4

Physical Adapters

- vmnic2 1000 Full

Advanced Switch Settings - Networking

- ESX Server 3.x provides some more sophisticated network settings

The screenshot shows the 'NIC Teaming' tab in the ESX Server 3.x configuration interface. It features several settings for network failover and load balancing, each with a checked checkbox and a dropdown menu. Below these settings is a section for 'Failover Order' with an 'Override vSwitch failover order' checkbox. At the bottom, there is a table listing network adapters, categorized into Active, Standby, and Unused. To the right of the table are 'Move Up' and 'Move Down' buttons.

General | Security | Traffic Shaping | **NIC Teaming**

Policy Exceptions

Load Balancing: Route based on the originating virtual port ID

Network Failover Detection: Link Status only

Notify Switches: Yes

Rolling Failover: No

Failover Order:

Override vSwitch failover order:

Select active and standby adapters for this port group. In a failover situation, standby adapters activate in the order specified below.

Name	Speed	Networks
Active Adapters		
vmnic1	1000 Full	192.168.52.1-192.168.52.254
vmnic2	1000 Full	
Standby Adapters		
Unused Adapters		

Move Up

Move Down

Using Replication Monitor

■ Validating Inbound Connections

The screenshot displays the Active Directory Replication Monitor application. The left pane shows a tree view of monitored servers under 'Default-First-Site-Name', including 'dc1', 'dc2', 'child1', and 'child2'. Each server has several replication objects listed, such as 'DC=classroom,DC=local' and 'CN=Configuration,DC=classroom,DC=local'. The right pane shows the 'Server Properties' dialog for a selected server, with the 'Inbound Replication Connections' tab active. It displays a table of inbound connection objects and a detailed description of one connection.

Monitored Servers:

Update Automatically

Log: C:\Documents and Settings\Administrator\My Documents\Default-First-Site-Name-
Status as of: 9/3/2006 12:12:07 PM

Server Properties

General | Server Flags | FSMO Roles | Credentials |

Inbound Replication Connections | TCP/IP Configuration

Below is a list of the inbound connection objects established for this domain controller:

Connection Name	Server Name	Admin-Generated?
0538dddd-32ad-49cd-9...	Default-First-Site-Na...	AUTO
360e077a-2da0-4043-8...	Default-First-Site-Na...	AUTO
ef98af74-fcb6-4738-a2...	Default-First-Site-Na...	AUTO

Below is an advanced description of why the inbound connection objects were established:

Connection Name: 0538dddd-32ad-49cd-92e9-160bccf56312

Replication Partner: Default-First-Site-Name\DC2
Administrator Generated?: AUTO

Reasons for this connection:

Directory Partition (DC=DomainDnsZones,DC=classroom,DC=local)
Replicated because the replication partner is a ring neighbor.

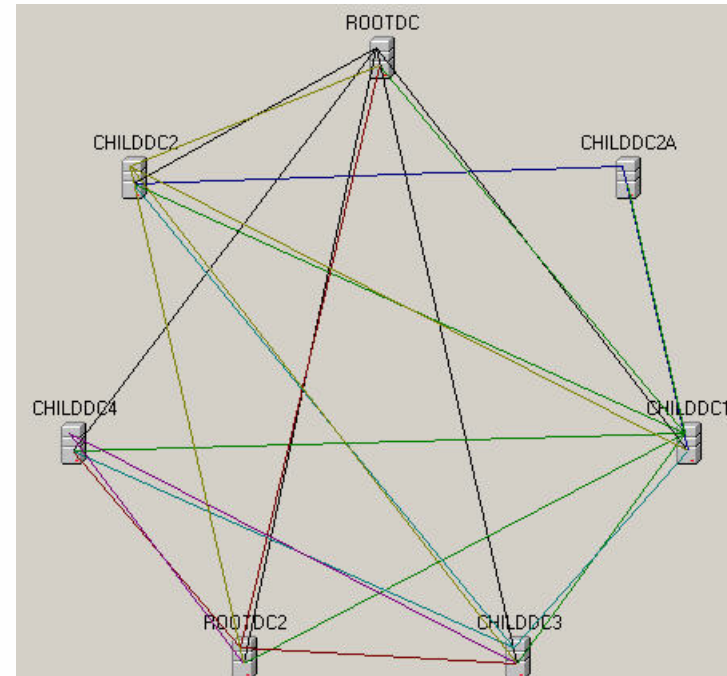
Using Replication Monitor (2)

■ Successful Replication

The screenshot displays the Active Directory Replication Monitor window. The left pane shows a tree view of monitored servers under 'Default-First-Site-Name', including 'dc1', 'dc2', 'child1', and 'child2'. Each server has a sub-tree of replicated objects such as 'DC=classroom,DC=local', 'CN=Configuration,DC=classroom,DC=local', 'CN=Schema,CN=Configuration,DC=classroom,DC=local', 'DC=DomainDnsZones,DC=classroom,DC=local', 'DC=ForestDnsZones,DC=classroom,DC=local', and 'DC=childb,DC=classroom,DC=local'. The right pane shows the log file path 'C:\Documents and Settings\Administrator\My Documents\dc1-DC=classroom,DC=local-DC2.log' and the status as of '9/3/2006 12:12:07 PM'. The log content indicates successful replication: '>> Direct Replication Partner Data <<' followed by 'Server is current through Property Update USN: 16594' and 'The last replication attempt was successful. This took place at: 9/3/2006 12:07:28 PM (local)'.

Replication Topology

- Checking Replication Topology



- Look for replication errors

<p>Monitored Servers</p> <ul style="list-style-type: none">Default-First-Site-Name<ul style="list-style-type: none">dc1<ul style="list-style-type: none">DC=classroom,DC=local<ul style="list-style-type: none">Default-First-Site-Name\DC2CN=Configuration,DC=classroom,DC=local<ul style="list-style-type: none">Default-First-Site-Name\DC2Default-First-Site-Name\CHILD2CN=Schema,CN=Configuration,DC=classroom,DC=local<ul style="list-style-type: none">Default-First-Site-Name\DC2Default-First-Site-Name\CHILD2	<p>g: C:\Documents and Settings\Administrator\My Documents\dc1-CN=Schema,CN=Configuration,DC=</p> <p>Status as of: 9/3/2006 12:12:07 PM</p> <p>>> Direct Replication Partner Data <<</p> <p>Server is current through Property Update USN: 5545</p> <p>Replication Failure: Changes have not been successfully replicated from DC2 for 1 attempt(s).</p> <p>Replication Failure: The reason is: There are no more endpoints available from the endpoint map</p> <p>Replication Failure: The last replication attempt was: 9/3/2006 11:48:42 AM (local)</p>
---	--

DNS Modifications

- Modify the weight and/or priority of the DNS SRV records
- Specifically offload the authentication requests from the PDC emulator when possible
- DNS **weight** is the proportional distribution of requests among DNS servers
- DNS **priority** is the likelihood a server will receive a request
- PDC emulators should have one or both adjusted accordingly by adding:
HKLM\System\CurrentControlSet\Services\Netlogon\Parameters
 - **LdapSrvWeight** **DWORD** decimal value of **25** or **50**
- PDC emulators should have one or both adjusted accordingly by adding:
HKLM\System\CurrentControlSet\Services\Netlogon\Parameters
 - **LdapSrvPriority** **DWORD** decimal value to **100** or **200**
- Physical domain controllers should be adjusted similarly to PDC emulator to decrease DNS dependencies on them

DNS Modifications

- Can also be changed within DNS manager
- Registry changes do not require a reboot

Service Location (SRV) | Security

Domain: Default-First-Site-Name.

Service: _ldap


Protocol: _tcp

Priority: 200

Weight: 50

Port number: 389

Host offering this service: rootdc.classroom.local.



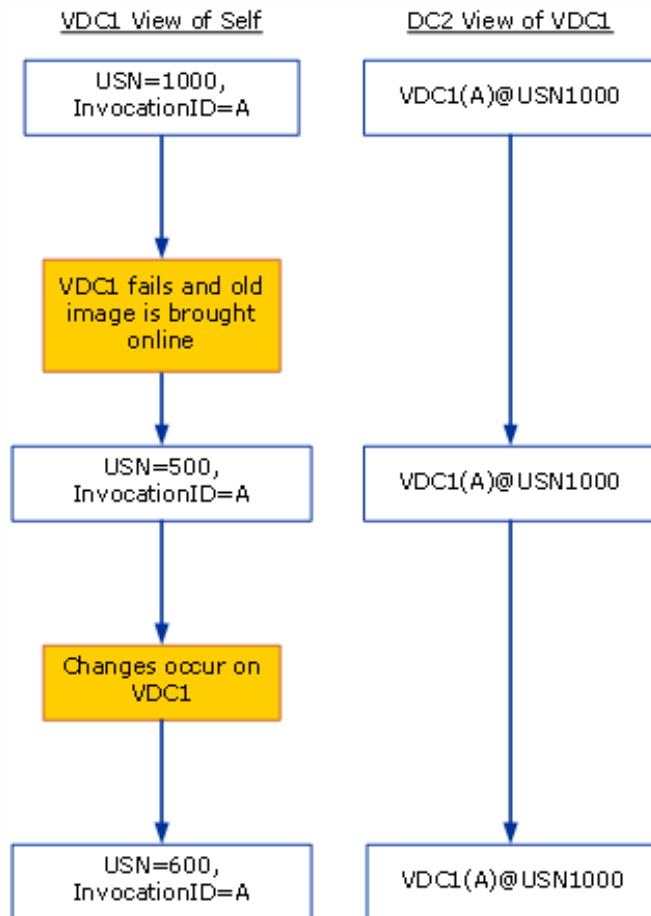
Name	Type	Data
_kerberos	Service Location (SRV)	[200][50][88] rootdc.classroom.local.
_ldap	Service Location (SRV)	[200][50][389] rootdc.classroom.local.
_gc	Service Location (SRV)	[200][50][3268] rootdc.classroom.local.
_kerberos	Service Location (SRV)	[0][100][88] rootdc2.classroom.local.
_ldap	Service Location (SRV)	[0][100][389] rootdc2.classroom.local.

Disaster Recovery

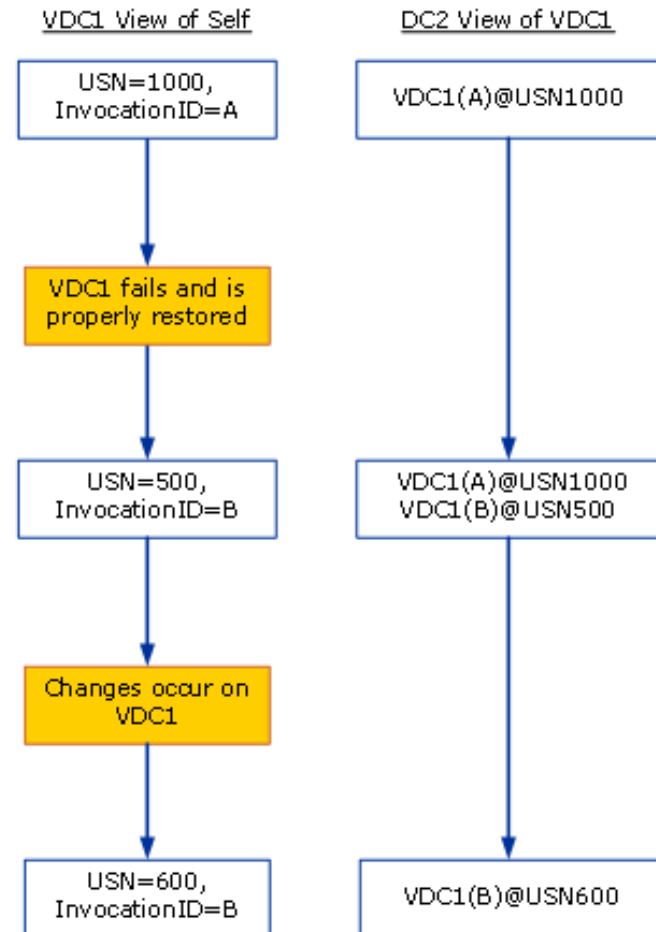
- Perform consistent system state backups
 - Eliminates hardware incapability when performing restore
- Follow Microsoft recommendations on FSMO role placement
 - <http://support.microsoft.com/kb/223346>
- All Active Directory restorations should be performed using authoritative and non-authoritative technique
 - Do not recover an Active Directory database from a backup copy of an old virtual disk!

Disaster Recovery Scenarios

Improper Restore of VM



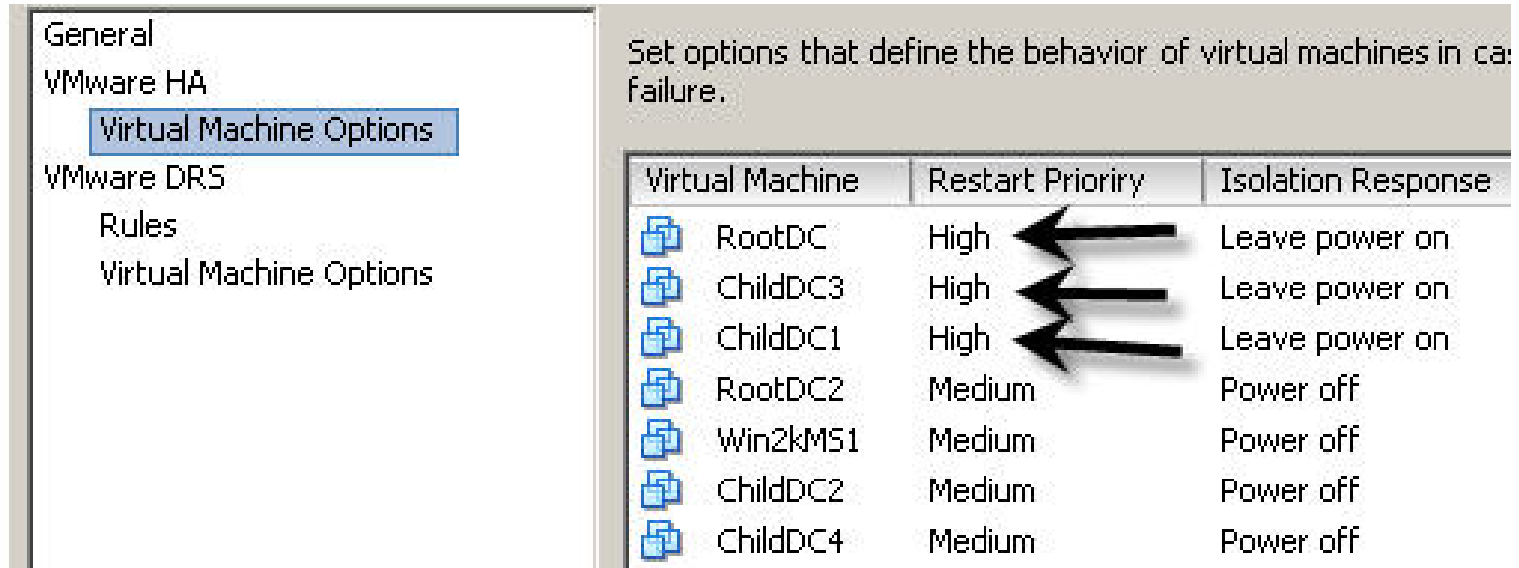
Proper Restore of VM



Source: Microsoft Corporation








Disaster Preparedness – ESX 3.x/VirtualCenter 2.x

- VMware provides solutions for automatically restarting virtual machines
- Implement VMware HA as a high availability to ensure virtual machine domain controllers restart in the event an ESX server fails



General
VMware HA
Virtual Machine Options
VMware DRS
Rules
Virtual Machine Options

Set options that define the behavior of virtual machines in case of failure.

Virtual Machine	Restart Priority	Isolation Response
 RootDC	High	Leave power on
 ChildDC3	High	Leave power on
 ChildDC1	High	Leave power on
 RootDC2	Medium	Power off
 Win2kMS1	Medium	Power off
 ChildDC2	Medium	Power off
 ChildDC4	Medium	Power off

Disaster Preparedness – ESX 3.x/VirtualCenter 2.x

- Combined with VMware DRS Anti-affinity rules can ensure domain controller VMs are segregated

The screenshot displays the VMware vSphere configuration interface for creating a DRS Anti-Affinity rule. The interface is divided into three main sections:

- Left Panel (Navigation):** Shows the 'General' tab selected under 'VMware DRS', with 'Rules' highlighted.
- Center Panel (Rule Selection):** Contains the text 'Use this page to create and apply to virtual machines. The rule will be retained if the virtual machine is moved.' Below this is a tree view of virtual machines with checkboxes and icons. The 'Root DCs' group is selected, and its sub-items 'RootDC' and 'RootDC2' are also selected.
- Right Panel (Rule Configuration):** Contains the following fields:
 - Name:** 'Root DCs'
 - Type:** 'Separate Virtual Machines' (indicated by an arrow from the 'Root DCs' group in the center panel)
 - Virtual Machines:** A list containing 'RootDC' and 'RootDC2'

Additional Information

- VMware Time Sync and Windows Time Service
 - VMware Knowledge Base ID# 1318
- Installing and Configuring NTP on VMware ESX Server
 - VMware Knowledge Base ID# 1339
- VMware Descheduled Time Accounting
 - http://www.vmware.com/pdf/vi3_esx_vmdesched.pdf
- How to detect and recover from a USN rollback in Windows Server 2003
 - <http://support.microsoft.com/kb/875495>
- How to detect and recover from a USN rollback in Windows 2000 Server
 - <http://support.microsoft.com/kb/885875>
- Support policy for Microsoft software running in non-Microsoft hardware virtualization software
 - <http://support.microsoft.com/kb/897615>
- How to configure an authoritative time server in Windows Server 2003
 - <http://support.microsoft.com/kb/816042>

Best Practices

- Avoid snapshots or REDOs for domain controller virtual machines
- Do not suspend domain controller virtual machines for long periods
- Consistent and regular system state backups still very important

Summary

- System State backups regularly
- Time Synchronization
- Disaster Recovery Plan
- High Availability
- Monitor Replication Traffic
- Modify DNS SRV records to redirect log on authentications to VMs
- Go back and constantly re-evaluate your strategy!!!

Thank you!!



VMWORLD 2006

Questions?



VMWORLD 2006