CONTINUITY SOFTWARE

vSphere Resilience HealthCheck

Free vCenter Analysis Report

User Name
Company Name

November 2, 2015

Powered by
AvailabilityGuard™

Introduction

This is an analysis of your VMware private cloud environment powered by AvailabilityGuard™. The report provides important insights on your VMware-based cloud infrastructure resiliency and measures your IT infrastructure's quality and data risk levels compared to other similar environments within your industry. The free report contains partial findings. Please contact us to get the complete report.

Detected vs. Inspected

The scope of the scan indicates which components were detected in the environment and which of them were inspected by the risk-detection engine.

VIRTUAL INFRASTRUCTURE	Detect	Inspect
vCenter servers	1	1
Virtual datacenters	3	3
ESX clusters	12	12
ESXi hosts	178	178

VIRTUAL MACHINES	Detect	Inspect
Windows	1370	×
Linux	415	×
Solaris	4	×

STORAGE	Detect	Inspect
Symmetrix arrays	8	×
NetApp Filers	4	×
Datastores	216	216

OTHER COMPONENTS	Detect	Inspect
Databases	×	×
App Servers	×	×
SRM protection plans	14	×

× Not included in the free edition

Finding Analysis

OVERALL INFRASTRUCTURE RISK LEVEL

TOTAL RISKS DETECTED » 58

RISK BY SEVERITY

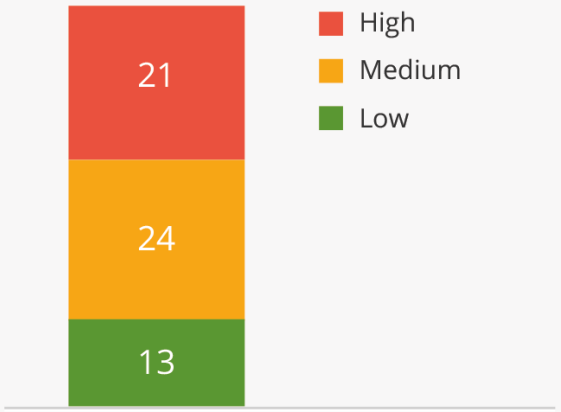

RISK BY IMPACT CATEGORY

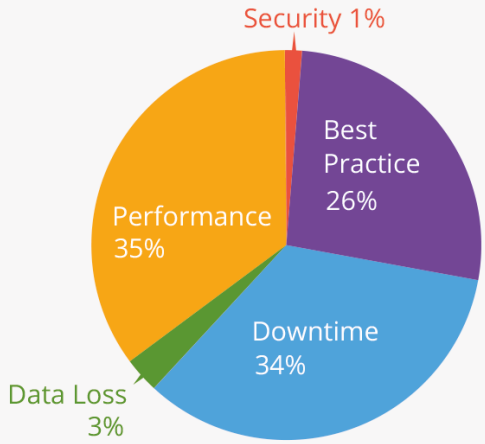

Risk by Impact Category

The following table summarizes risks by category, number of issues, impact and suggested priority.

CATEGORY	# RISKS	INCLUDED	IMPACT	SEVERITY
ESX cluster has incorrect timekeeping (NTP) configuration	3	1	Downtime	High
Transparent Page Sharing (TPS) is enabled on cluster	1	1	Security	High
VM port group configuration inconsistency was detected between hosts of ESX cluster	7	1	Downtime	Medium
Datastores of ESX cluster are not configured on all cluster nodes	2	1	Downtime, Best Practice	High
Mounted CD-ROM on an HA protected VM	1	1	Downtime, Performance	High
HA and vMotion clusters not configured correctly with EVC	2	1	Downtime	Medium
SRM manages ESX servers of unrecommended version	2	1	Downtime	Medium
SRM License over-allocated	2	1	Downtime	Medium
Suboptimal NFS MaxQueueDepth	2	1	Performance	Medium
Hosts accessing shared storage with different SAN I/O configuration	3	1	Performance	Low
CPU Hotplug is enabled on vNUMA VMs	3	1	Performance	Medium
Resource allocation limits should not be used for VMs	4	1	Performance, Best Practice	Medium
ESX cluster has physical volumes managed by VMware MPIO which has inconsistent LUN number	2	0	Data Loss, Best Practice	Medium
Changed Block Tracking (CBT) is disabled for 6 virtual machines	1	0	Performance, Best Practice	Low
ESX hosts using LUNs with dead paths	6	0	Downtime	High
11 more categories	17	0		
26 categories detected.	**58**	**12**	Click here to learn more	

Detailed Risk Information

For each detected risk, AvailabilityGuard creates a detailed ticket to capture the risk's contextual-information, explain possible impact, suggest remediation steps and provide optional diagrams and user notes.

Risk 622

Name	Incorrect NTP configuration for ESX host
Severity	High
Categories	Best Practice

Summary

ESX cluster **moon** at sites **Chicago** and **Boston** has incorrect timekeeping (NTP) configuration.

Description

A gap has been detected where 2 out of 4 hosts of ESX cluster moon do not synchronize time from an NTP server. The VMware best practice is to configure an authoritative time (NTP) server for ESX/ESXi hosts.
This issue is caused by one or more of the following reasons:

- NTP server is not configured for the host.
- The NTP client service is disabled.
- Ports required for NTP communication are blocked by the host firewall.

Incorrect NTP configuration may lead to inaccurate time on ESX hosts and for virtual machines that rely on their host for time synchronization. Consequentially it may jeopardize the stability and availability of the ESX hosts and virtual machines (see impact).

The following table presents key NTP configuration information for all the hosts of ESX cluster moon

Host name	Defined NTP servers	NTP client traffic allowed in the host firewall	NTP service daemon status
moon2	time.microsoft.com	Yes	Running
moon1	Undefined	Yes	Running
moon3	Undefined	Yes	Running
moon4	time.microsoft.com	Yes	Running

No virtual machines are currently configured to synchronize time with their underlying host.

Impact

The internal ESX/ESXi host clock may drift from the real time and, as a result, the host will suffer from an inaccurate time.
Time-sensitive services and applications running on the host and on the virtual machines may suffer from unexpected errors. Such unexpected errors may include authentication issues for hosts and virtual machines (when AD/Kerberos used), incorrect timestamp logged for transactions in databases, and more.
In addition, external management systems that communicate with the unsynchronized host or virtual machines may be affected as well. Last, various services require NTP configuration for ESX/ESXi servers for correct deployment and operations (VMware vFabric Data Director, Cisco Virtual Network Management Center, Avaya Secure Access Link, and more).

Resolution

Configure a reliable NTP server on your ESX hosts. The authoritative time server could be a Microsoft Active Directory Server or Internet time server.
If an Internet time server is used, take care to make the necessary changes to the corporate firewall to enable NTP communication.
In general, it is advised to configure multiple NTP servers and use a local server as the primary NTP server.

Topology

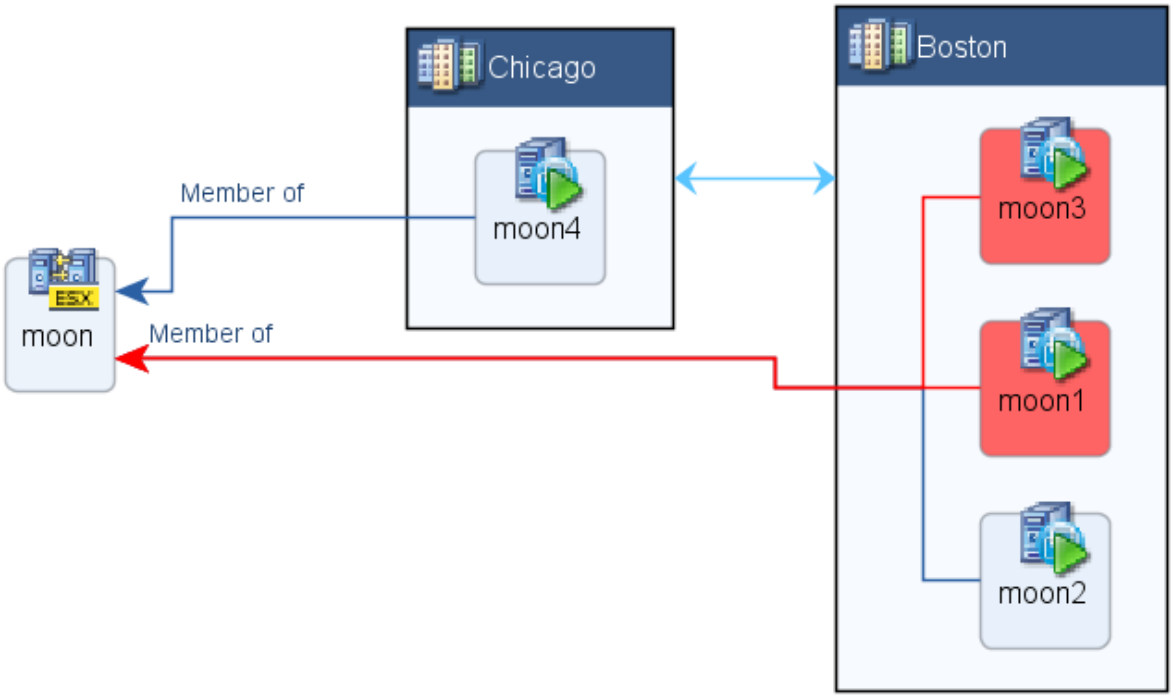
Chicago
Boston
Member of
moon4
moon3
moon
Member of
moon1
moon2

Risk 507

Name	Transparent Page Sharing is enabled
Severity	High
Categories	Best Practice

Summary

Transparent Page Sharing (TPS) is enabled on ESX cluster **NASA** at site **Chicago**.

Description

A gap was found where hosts of ESX cluster **NASA** at site **Chicago** are configured with the TPS setting enabled.
This configurations puts virtual machines running on ESX cluster **NASA** at site **Chicago** at the risk of unauthorized access and breach of AES encryption (see impact)

TPS is enabled on the following hosts:

Host name	Number of hosted VMs
host851	12
host853	10
host852	8

Impact

By forcing a flush and reload of cache memory, it is possible to measure memory timings to assist in dramatically accelerating the time it takes to determine an AES encryption key in use on another virtual machine running on the same physical processor of the host server (certain theoretical papers suggest the exploit could be extended to additional hosts in the cluster as well.)
This attack could be performed on any other VM that also has TPS enabled running on the same host (and possibly, the same cluster).

Resolution

To disable TPS for ESXi 5.x, perform the following steps:

- Log in to ESX\ESXi or vCenter Server using the vSphere Client. If connected to vCenter Server, select the relevant ESX\ESXi host.
- In the Configuration tab, click Advanced Settings under the software section.
- In the Advanced Settings window, click Mem.
- Look for Mem.ShareScanGHz and set the value to 0.
- Click OK.

Topology

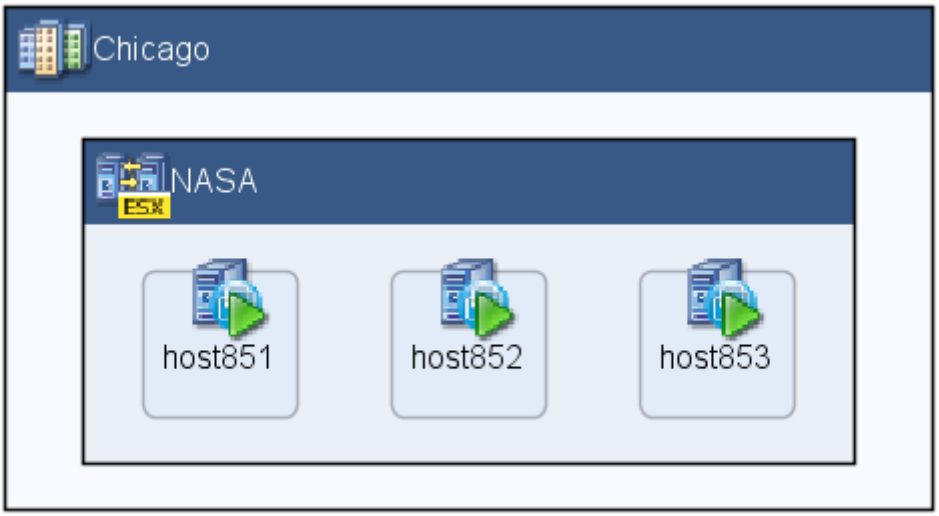

Risk 540

Name	VM port group configuration inconsistency between cluster hosts
Severity	Medium
Categories	Downtime

Summary

VM port group configuration inconsistency was detected between hosts of ESX cluster **Cluster B** at site **Chicago**.

Description

A gap has been detected where ESX cluster **Cluster B** at site **Chicago** has inconsistent VM port group definitions. Certain VM port groups are defined on some of the cluster hosts but not on all of them. This configuration may adversely impact the availability of virtual machines and successful completion of common VMware processes such as VM restart on another cluster host (see impact).

The following table identifies the VM port groups of cluster Cluster B and the hosts on which they are defined:

vSwitch	Port Group Name	Defined on Hosts	Not Defined on Hosts
vSwitch0	VM Network	host411, host412, host414, host415, host416, host417, host418	host413

The following table identifies the virtual machines dependent on missing port groups:

Host name	Missing Port Group	VMs assigned to this Port Group	Operating System
host411	VM Network	* vmdb30 * vmdb57	Linux
host412	VM Network	* vmdb14 * vmdb06	Solaris
host414	VM Network	* vmwindows10 * vmwindows03	Windows
host415	VM Network	* vmwindows07 * vmwindows05	Windows
host416	VM Network	* vmdb02 * vmdb09	Solaris

Impact

This configuration issue poses a downtime risk to virtual machines. When a VM will be restarted through VMware HA or DRS (or other mean) on a host that is incorrectly defined with its VM port groups, the VM will fail to gain access to all required networks. VMware expects consistent port group names on all the cluster hosts.

Resolution

Reconfigure VM port groups as needed. Consider leveraging VMware vSphere Distributed Switch (vDS) to eliminate network setup inconsistencies between hosts. Note that port group names are case sensitive.

Topology

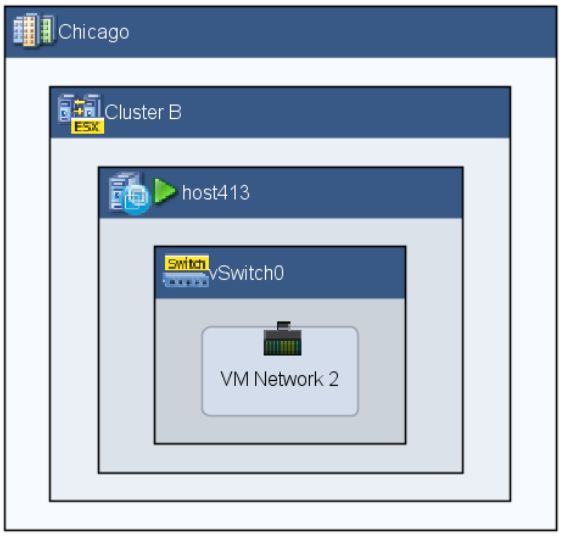

Risk 479

Name	Datastore not configured on all ESX cluster nodes
Severity	High
Categories	Downtime, Extended Recovery Time

Summary

Datastores of ESX cluster **MainProd-DR** at site **Chicago** are not configured on all cluster nodes.

Description

Datastores of ESX cluster **MainProd-DR** are not configured on all cluster nodes. This might result in prolonged downtime of dependent virtual machines (see impact).

The following table lists the datastores which are not configured in all nodes of ESX cluster MainProd-DR:

Datastore	Configured on hosts	Not configured on hosts
Sharepoint-DS01	* host386 * host387 * host388 * host389 * host390 * host391	host392

The following table lists DAS datastores used by ESX cluster MainProd-DR:

Datastore	Configured on hosts
Sharepoint-DS12	host386
Sharepoint-DS06	host387
Sharepoint-DS04	host388
Sharepoint-DS14	host389
Sharepoint-DS21	host390
Sharepoint-DS16	host391
Sharepoint-DS01	host392

The well-configured shared datastores used by ESX cluster **MainProd-DR**:

- **Sharepoint-DS12**
- **Sharepoint-DS06**
- **Sharepoint-DS04**
- **Sharepoint-DS14**
- **Sharepoint-DS21**
- **Sharepoint-DS16**

The virtual machines stored on **TestDR-DS01**:

- Windows VM **vm1194**
- Linux VM **vm343-main**
- Linux VM **vm428**
- Windows VM **vm1192**

Impact

Virtual machines fail-over to an alternate cluster node is possible only if that node has access to all datastores the virtual machines rely on. Since some of the cluster nodes have only partial access to the specified datastores the dependent virtual machines might suffer an extended downtime upon cluster node outage.

Resolution

Make sure that the underlying datastores storage volumes are mapped to all ESX nodes. Rescan storage on the cluster nodes that do not use the specified datastores.

Topology

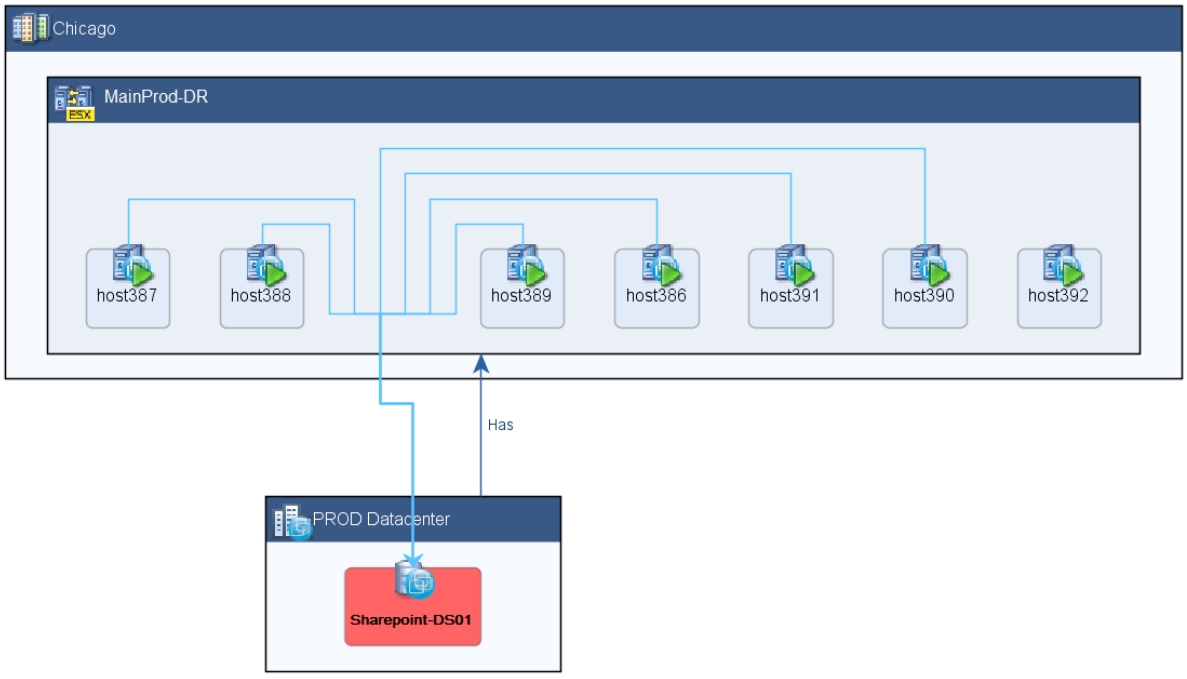

Risk 480

Name	Mounted CD-ROM on an HA protected VM
Severity	High
Categories	Availability

Summary

Virtual machines running on ESX cluster **Cluster B** at site **Chicago** mount a CD-ROM are dependent on local storage that is inaccessible to other cluster hosts.

Description

A gap has been detected where virtual machines in ESX cluster **Cluster B** at site **Chicago** mount a CD-ROM dependant on local storage that is inaccessible to other cluster hosts. This configuration will lead to unsuccessful vMotion (see impact).

The following virtual machines have a mounted CD-ROM on storage local to the host:

VM name	OS Type	CD-ROM Connected	CD-ROM Connected On Powered-On	Running	Datastore file
ui vm	Linux	No	Yes	No	ISO [Host412.datastore] indeni.4.0.6.iso

Impact

In the event of vMotion the VM will fail to load on other cluster hosts because the CD-ROM device is mounted on a datastore which is inaccessible to them.

Resolution

Consider un-mounting the CD-ROM device on the virtual machines.

Topology

Risk 6809

Name	EVC is not used in some clusters
Severity	Medium
Categories	Best Practice, Availability

Summary

ESX cluster **MilkyWay** that has both vMotion and HA enabled and EVC mode incorrectly configured.

Description

A gap has been detected where ESX cluster **MilkyWay** that has both vMotion and HA enabled and EVC mode incorrectly configured. The lack of EVC mode configuration might result in unexpected downtime and unbalanced hosts (see impact).

The following HA-enabled hosts are not configured with EVC mode:

Max EVC Mode	Hosts
intel-sandybridge	HostDB010, HostDB011, HostApp012, HostApp013
intel-westmere	HostApp010, HostApp011, HostDB012, HostDB013, HostDB015

Impact

On a mixed cluster if EVC mode is not used, the following can occur:

- VMs will be moved using vMotion in one direction only (from hosts with a less current architecture to those with more current one). This could lead to an un-balanced cluster, and will disrupt DRS' load balancing algorithms
- Certain non-well behaved applications might cause VMs to crash when installed on a host with a more current application and then restarted (with or without VMware HA) on a host with a less current one.

Resolution

Consider configuring the EVC mode on the cluster.

Risk 554

Name	SRM manages ESX servers of unrecommended version
Severity	Medium
Categories	Downtime, Best Practice

Summary

VMware SRM **SRM10** is managing ESX servers that have an earlier and un-recommended version.

Description

A gap has been detected where the SRM server **SRM10** is managing ESX servers of version earlier than 5.0. This might lead to a Permanent Device Loss (PDL) (see impact).

The following table lists the ESXs servers with earlier versions:

ESX name	Site	Type	Version
moon3	Boston	ESXi	4.0.0
moon4	Chicago	ESXi	4.0.0
moon1	Boston	ESXi	4.0.0
moon2	Boston	ESXi	4.0.0

Impact

Using SRM 5.0 and above with ESX of versions earlier than 5.0 may result in a state called Permanent Device Loss ("PDL"). PDL may occur during planned migrations and SRM test failovers.
According to VMware: "When SRM is protecting virtual machines running on ESX 5.0 hosts, many cases where PDL could occur are properly handled, and PDL is avoided".
PDL (as defined by VMWare):
Occurs when the storage array returns SCSI sense codes indicating that the LUN is no longer available or that a severe, unrecoverable hardware problem exist with it.

Resolution

Upgrade the ESX servers to version 5 or higher.

Topology

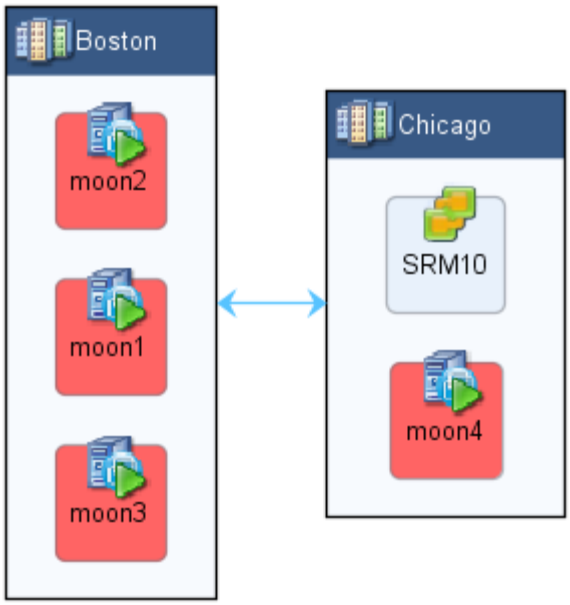

Risk 127

Name	SRM License over-allocated
Severity	Medium
Categories	Best Practice

Summary

The SRM License 'RJJTA-#####-#####-####C-2OSKD3' at site '**Boston**' is over-allocated by 7 licenses.

Description

A gap has been detected where the SRM license currently allocated to the **Boston** vCenter

Impact

Upon learning that product licenses are over-allocated, SRM will constantly generate events to the vCenter console, urging admins to renew their licenses. In the case of evaluation licenses, SRM will cease working once its licenses are over-allocated.

Resolution

SRM licenses are required only for the hosts at the protected site that are running the protected virtual machines. Protected VMs turned into unprotected VMs will reduce the amount of licenses used.

Risk 632

Name	Suboptimal NFS MaxQueueDepth
Severity	Medium
Categories	Downtime

Summary

Suboptimal NFS.MaxQueueDepth configuration on ESXi hosts on ESX cluster **Cluster B** at site **Chicago**.

Description

A gap has been detected where ESXi servers are configured with an exceptionally high and un-recommended value for the NFS max queue depth option: 4294967295. The recommended value is 64. This configuration may lead to serious NFS connectivity issues (see impact).

The following table lists the servers and their NFS max queue depth configuration:

Host	Datastores	VMs on effected datastores	Number of effected VMs	NFS MaxQueueDepth
host413	* vCenter_Prod (1) * vCenter_ Prod	* vm434 * vm1412 * vm435	7	4294967295
host412	* vCenter_ Prod (1) * vCenter_ Prod	* vm004 * ui vm	17	4294967295

Impact

- Affected NFS datastores appear to be unavailable (greyed out) in vCenter Server, or when accessed through the vSphere Client.
- The NFS intermittently disappear and reappear every few minutes.
- Virtual machines located on the NFS datastore are in a hung/paused state when the NFS datastore is unavailable.
- This issue is most often seen after a host upgrade to ESXi 5.x or the addition of an ESXi 5.x host to the environment.

Resolution

To set the NFS.MaxQueueDepth advanced parameter using the vSphere Client: Click the host in the Hosts and Clusters view.

- Click the Configuration tab, then click Advanced Settings under Software.
- Click NFS, then scroll down to NFS.MaxQueueDepth.
- Change the value to 64.
- Click OK.

Reboot the host for the change to take effect.

Topology

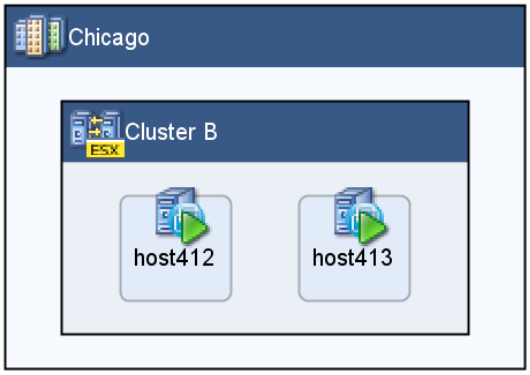

Risk 585

Name	Hosts accessing shared storage with different SAN I/O configuration
Severity	Low
Categories	Availability, Performance

Summary

Number of SAN I/O paths inconsistency between nodes of ESX cluster **Zoo** at site **Chicago**.

Description

Shared HDS logical units mapped to 3 nodes of ESX cluster **Zoo** are accessed with different number of SAN I/O paths. This might result in performance degradation (see impact).

The following section lists the inconsistent SAN I/O access configuration:

Data	Primary HDS logical unit	NMP - 2 paths (VMW_SATP_DEFAULT_AA / VMW_PSP_RR)	NMP - 4 paths (VMW_SATP_DEFAULT_AA / VMW_PSP_RR)
Datastore VSP_LOCAL2-DC-INT-Mng_01 of datacenter PROD Datacenter	12345/1401	* Zoo/host1384	* Zoo/host1383 * Zoo/host1382
Datastore VSP_LOCAL2-DC-INT-Mng_02 of datacenter PROD Datacenter	12345/1400	* Zoo/host1384	* Zoo/host1383 * Zoo/host1382
Datastore VSP_LOCAL2-DC-INT-Mng_03 of datacenter PROD Datacenter	12345/1411	* Zoo/host1384	* Zoo/host1383 * Zoo/host1382
Datastore zScratch_LOCAL2-DC-Mgn_VSP of datacenter PROD Datacenter	12345/FC37	* Zoo/host1384	* Zoo/host1383 * Zoo/host1382

Impact

Using different multipathing solutions might affect the service level experienced by virtual machines that use the identified storage volumes. When a virtual machine is relocated as a result of a DRS or HA event, it may demonstrate a change I/O performance.

Resolution

- Determine how many paths are required for the SAN storage volumes, what is the required I/O policy, and if relevant, the storage array type plugin
- Reconfigure the devices with inadequate configuration according to those requirements

Strive to ensure end-to-end path redundancy (e.g., using separate array ports, fabric switches and host HBAs) as much as possible

Topology

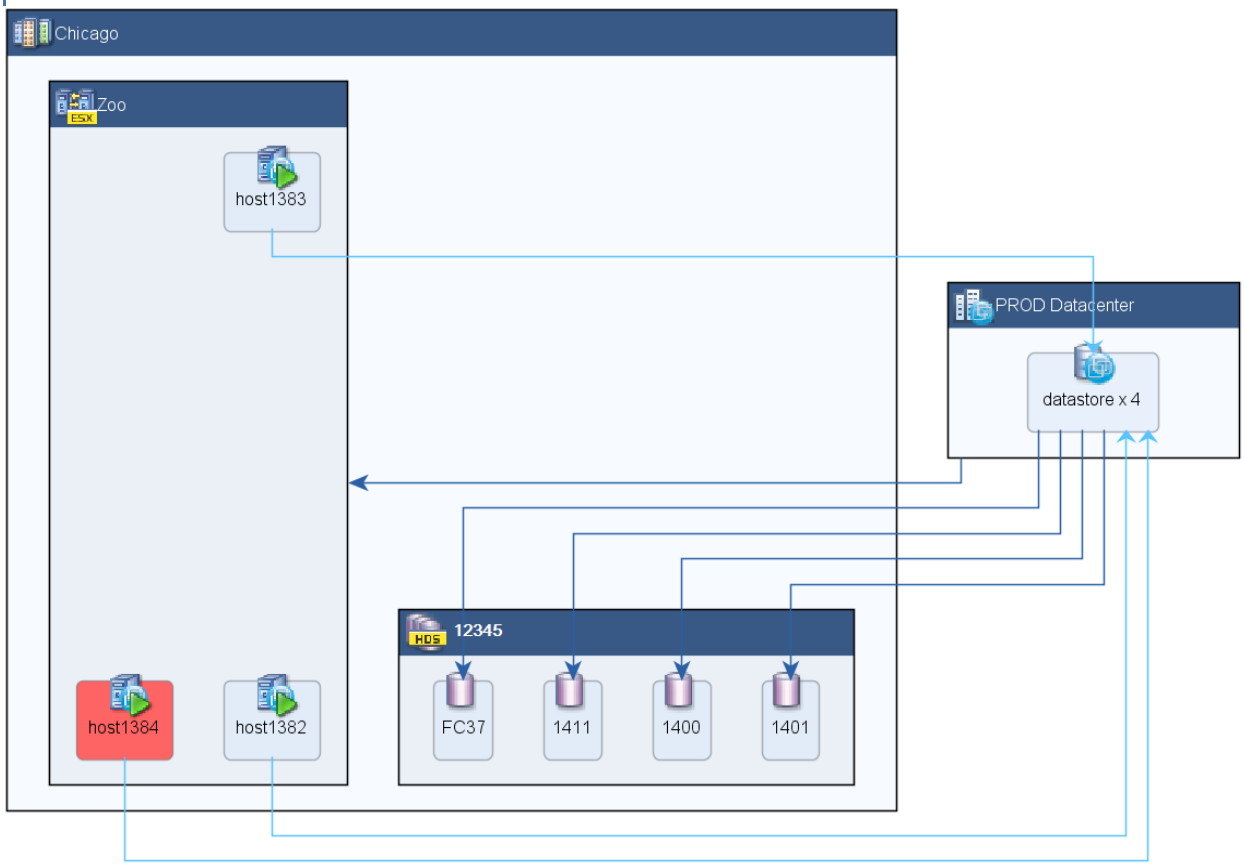
Chicago
Zoo
ESX
host1383
host1384
host1382
PROD Datacenter
datastore x 4
12345
HDS
FC37
1411
1400
1401

Risk 527

Name	CPU Hotplug is enabled on vNUMA VMs
Severity	Medium
Categories	Performance

Summary

CPU Hotplug is enabled on vNUMA-configured VMs in ESX cluster **MilkyWay** at site **Chicago**.

Description

A gap has been detected where CPU Hotplug is enabled on vNUMA-configured VMs in ESX cluster **MilkyWay** at site **Chicago** resulting in vNUMA being effectively disabled. This might lead to performance degradation (see impact).

The following vNUMA-enabled virtual machines also have CPU Hotplug enabled:

VM name	OS running on VM	Configured memory	Configured vCPU
dwh01	Microsoft Windows Server 2012 R1 (64bit)	32768	4
dwh02	Microsoft Windows Server 2012 R1 (64bit)	32768	4
dwh03	Microsoft Windows Server 2012 R1 (64bit)	32768	4
dwh04	Microsoft Windows Server 2012 R1 (64bit)	32768	4
dwh05	Microsoft Windows Server 2012 R1 (64bit)	32768	4

Impact

When CPU Hotplug is enabled, vNUMA is effectively disabled, and the VM will be started with Uniform Memory Access instead. This will decrease memory access performance and might significantly impact memory-intensive applications. Note that vNUMA is not enabled by default on VMs with less than 8 vCPUs.

Resolution

To disable CPU Hotplug:

- In the vSphere Client inventory, right-click the virtual machine and select Edit Settings.
- Click the Options tab and under Advanced, select Memory/CPU Hotplug.
- Change the CPU Hot Plug setting to disabled.
- Click OK to save your changes and close the dialog box.

Resolution

To disable CPU Hotplug:

- In the vSphere Client inventory, right-click the virtual machine and select Edit Settings.
- Click the Options tab and under Advanced, select Memory/CPU Hotplug.
- Change the CPU Hot Plug setting to disabled.
- Click OK to save your changes and close the dialog box.

RISK 6309

Name	Resource allocation limits should not be used for VMs
Severity	Medium
Categories	Best Practice

Summary

VMs of ESX cluster **cloudDB** are configured with resource allocation limits.

Description

A gap has been detected where on some VMs of cluster **cloudDB** resources limits are imposed, contrary to the best practice. This might result in performance degradation of the affected VMs (see impact).

The following VMs are configured with limits:

VM	VM Name	CPU Limits	Memory Limits
vmdb2-100	vmdb2-100	9572	12288
vmdb2-200	vmdb2-200	9572	12288

Impact

When a memory limit is set lower than the virtual machine's provisioned memory, it is considered the upper boundary for the amount of physical memory that can be directly assigned to this particular virtual machine. The guest operating system is not aware of this limit, and it optimizes memory management options to the assigned memory size.
When the limit is reached or exceeded, the guest operating system can still request new pages, but due to the limit the VMkernel does not allow the guest to directly consume more physical memory and treats the virtual machine as if the resource is under contention. As such, memory reclamation techniques are used to enable the virtual machine to consume what it has requested.

Resolution

Consider removing all limits from the affected VMs. Shares could be used to better restrict resource usage of certain virtual machines only when the host becomes overcommitted.

About the AvailabilityGuard Technology

AvailabilityGuard™ enables IT teams to proactively identify and eliminate misconfigurations and single-points-of-failure across the entire IT infrastructure — including High Availability, Cloud, and Disaster Recovery (DR) environments.

AvailabilityGuard's predictive IT Operations Analytics, helps top banks, major financial institutions and many other leading organizations to pinpoint potential failures before they impact the business — delivering the highest levels of IT service availability while improving operational efficiency.

Proactive Risk Detection

Using a predictive risk-detection engine, AvailabilityGuard verifies infrastructure resilience on a daily basis. It can automatically identify over 5,000 cross-vendor misconfigurations across all IT layers that may lead to outages, alerting the appropriate IT teams to take action and arming them with the suggested resolution.

- **Detect** issues that can affect service availability and data by performing non-intrusive, cross-domain routine scans and predictive analysis of your IT infrastructure configuration.
- **Alert** the appropriate teams and enable them to collaborate on a solution.
- **Correct** configuration issues before they impact the business and require costly firefighting.

The AvailabilityGuard Dashboard

The AvailabilityGuard dashboard provides immediate visibility into availability and data loss risks throughout your entire IT infrastructure and their potential impact on critical business services. From the dashboard, you are just a click away from detailed information on any issue.

Learn more about how AvailabilityGuard works and our Enterprise solution.