AvailabilityGuard cloud edition

Private Cloud Health Check Report

A total of 15 potential downtime and data loss risks were detected by AvailabilityGuard/Cloud health check assessment of your private cloud environment.

Issues are assigned a priority according to the following criteria:

- Priority 1: risk of extended downtime or data loss to a mission-critical business service
- Priority 2: risk of extended downtime or data loss in any production system
- Priority 3: risk of performance degradation in a production system
- Priority 4: any risk to non-production system
- Priority 5: deviation from best practices

HIGH PRIORITY THREATS

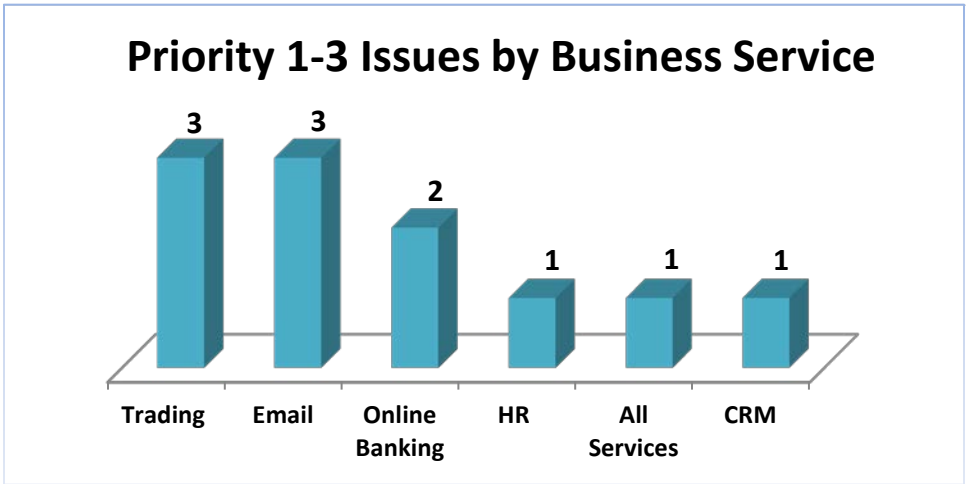

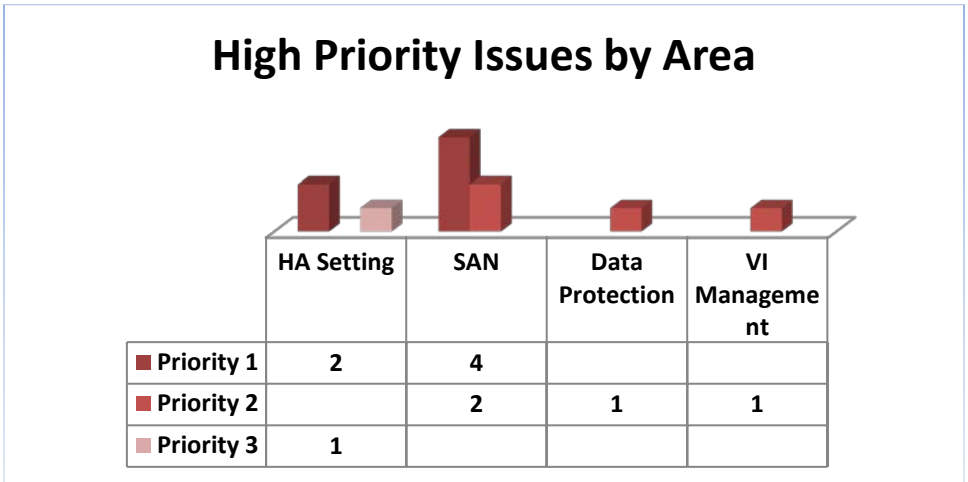

FINDING SUMMARY

Affected Business Services	Area	Category	Possible Impact	Priority	# of Open Tickets	Tickets
Online Banking	HA Setting	Cluster configuration errors	• Extended downtime (single point of failure in a system designed for High Availability) • Potential data loss upon fail-over	1	2	498, 40
Trading	SAN	SAN I/O configuration issues	• Single Point of Failure • Risk of data corruption	1	3	62, 28, 230
CRM	SAN	Unauthorized storage access	• Risk of massive data corruption affecting multiple Virtual Machines	1	1	31
All Services	VI Mgmt	VI best practices	• Inability to manage the VI	2	1	43
Email, DWH	SAN	Incorrect storage array configuration	• Reduced performance • Reduced service availability	2	2	46, 173
Email	Data Protection	Cluster storage access	• No compliance with data retention requirement • Insufficient performance and redundancy expected after fail-over	2	1	185
HR	HA Setting	VI best practices	• Reduced performance • Inefficient utilization of I/O resources	3	1	92
Document Mgmt	Optimization	Optimization	• Sub-optimal performance • Opportunity to eliminate storage hot-spots and bottlenecks	5	4	42, 211, 118, 36

TICKET 498

Gap ID	00451DTSMIS	**Name**	Datastore not configured on all ESX cluster nodes
Severity	INFO	**Status**	OPEN
Categories	Downtime, Extended Recovery Time	**Rating**	★★★★★
Detected on	Nov 22, 2011 1:56:37 AM		
Verified on	Dec 28, 2011 1:53:10 AM		

SUMMARY

The datastore MAFS-PRD02 of ESX cluster ma_prd_01 at site Boston is not configured on all cluster nodes.

DESCRIPTION

The datastore MAFS-PRD02 of ESX cluster ma_prd_01 is not configured on all cluster nodes. This might result in prolonged downtime of dependent virtual machines (see impact).

The following table lists the datastores inconsistently configured on the ESX cluster ma_prd_01:

Datastore	db01	db03	db02	db03	db02
MAFS-PRD10	Configured	Configured	Configured	Configured	Configured
MAFS-PRD11	Configured	Configured	Configured	Configured	Configured
MAFS-PRD02	Configured	Configured	Configured	**Not configured**	Configured
MAFS-PRD07	Configured	Configured	Configured	Configured	Configured

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 datastore MAFS-PRD02, the dependent virtual machines might suffer an extended downtime upon cluster node outage.

RESOLUTION

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

TICKET 40

Gap ID	00462VMTCV	**Name**	ESX servers of different CPU architecture sharing a datastore
Severity	WARNING	**Status**	OPEN
Categories	Downtime	**Rating**	
Detected on	Jan 23, 2012 12:11:39 PM		
Verified on	Jan 30, 2012 12:48:59 AM		

SUMMARY

ESX cluster **DELL g6** at site **Tokyo** has nodes with different CPU vendor.

DESCRIPTION

VMware HA is feature of an ESX cluster designed to provide transparent virtual machine fail-over in the event of a physical node outage. Hardware configuration does not need to be identical across all cluster nodes, and some nodes may have less CPUs and memory than others. However, it is a best practice that all nodes in the same cluster will have the same CPU architecture and vendor, to prevent unexpected system crash upon virtual machine failover or VMotion (see impact).

The following table lists the ESX servers of different CPU architectures:

ESX host	CPU type
TKesx19	Intel(R) Xeon(R) CPU E5335 @ 2.00GHz
TKesx21	Intel(R) Xeon(R) CPU E5430 @ 2.66GHz
TKesx24	Dual-Core AMD Opteron(tm) Processor 2218

IMPACT

Attempting to perform live migration (VMotion) of a virtual machine between two physical servers of different CPU architectures will fail, running the risk of unplanned outage of the virtual machine.

RESOLUTION

Make sure that the all cluster nodes have the same CPU architecture.

TICKET 62

Gap ID	00245DeadPath	**Name**	Devices with dead paths
Severity	WARNING	**Status**	OPEN
Categories	Availability	**Rating**	★★★★★
Detected on	Jan 27, 2012 12:25:03 PM		
Verified on	Jan 30, 2012 12:46:14 AM		

SUMMARY

ESX cluster node **TKvc02** of ESX cluster **O2** at site **Tokyo** has dead SAN I/O paths.

DESCRIPTION

ESX cluster node **TKvc02** of ESX cluster **O2** at site **Tokyo** is using SAN storage volumes. Some of the SAN I/O paths to these storage volumes are unavailable. This might result in reduced performance and redundancy.

The following table shows the unavailable paths:

Local file system	PV	Dead path	SAN device	Unavailable paths/Total # of paths
Lun1	**HITACHI Fibre Channel Disk (naa.FF0000)**	fc.df09b:3c-fc.ffffffffffffffff:f2-naa.FF0000	**4472 /0000**	**1/2**
Lun2	**HITACHI Fibre Channel Disk (naa.FF0001)**	fc.df09b:3c-fc.ffffffffffffffff:f2-naa.FF0001	**4472 /0001**	**1/2**
Lun3	**HITACHI Fibre Channel Disk (naa.FF0003)**	fc.df09b:3c-fc.ffffffffffffffff:f2-naa.FF0003	**4472 /0003**	**1/2**

IMPACT

Having only one live path to storage volumes might lead to:

- Redundancy loss, single-point-of-failure issue
- Increased risk of downtime
- Degraded performance

RESOLUTION

Reconfigure the lost paths using your multi-pathing software. If the issue cannot be resolved at the software level - check related hardware for possible faults. Examples for hardware configuration issues:

- Incorrect zoning or masking definitions
- Switch outage
- Physical link problems
- Excessively long paths (e.g., with more than 6 hops)
- Incorrect mapping definitions at the storage array
- Etc.

TICKET 31

Gap ID	00478VMDSUAC	**Name**	Unauthorized SAN access to VMware datastore
Severity	ERROR	**Status**	OPEN
Categories	Tampering	**Rating**	★★★★★
Detected on	Sep 19, 2011 2:30:14 AM		
Verified on	Sep 27, 2011 2:49:31 AM		

SUMMARY

VMware datastore **VD12** defined in datacenter **London** is accessed through the SAN by unauthorized hosts.

DESCRIPTION

VMware datastore **VD12** defined in datacenter **London** at site **London** is using Symmetrix device which is accessed by ESX hosts that are not members of the same datacenter. This might lead to data corruption (see impact).

The following table presents the Symmetrix device and the unauthorized hosts:

Symmetrix device	Host	Device
4127 /102E	ESX host **ld-esx1**	**EMC Fibre Channel Disk (naa.fe4344)**
4127 /102E	ESX host **ld-esx2**	**EMC Fibre Channel Disk (naa.fe4344)**

IMPACT

Symmetrix devices allocated for datastore use should be mapped only to ESX servers that are members of the same datacenter. This can be useful to allow migration of virtual machines using VMotion. However, it is highly dangerous to allow SAN access to the underlying Symmetrix devices by any other host, especially if one is not an ESX server. This might result in unexpected data corruption and data loss that will affect any Virtual Machine using the datastore - which might also lead to downtime.

RESOLUTION

Remove the SAN mapping of the Symmetrix device to the unauthorized hosts.

TOPOLOGY

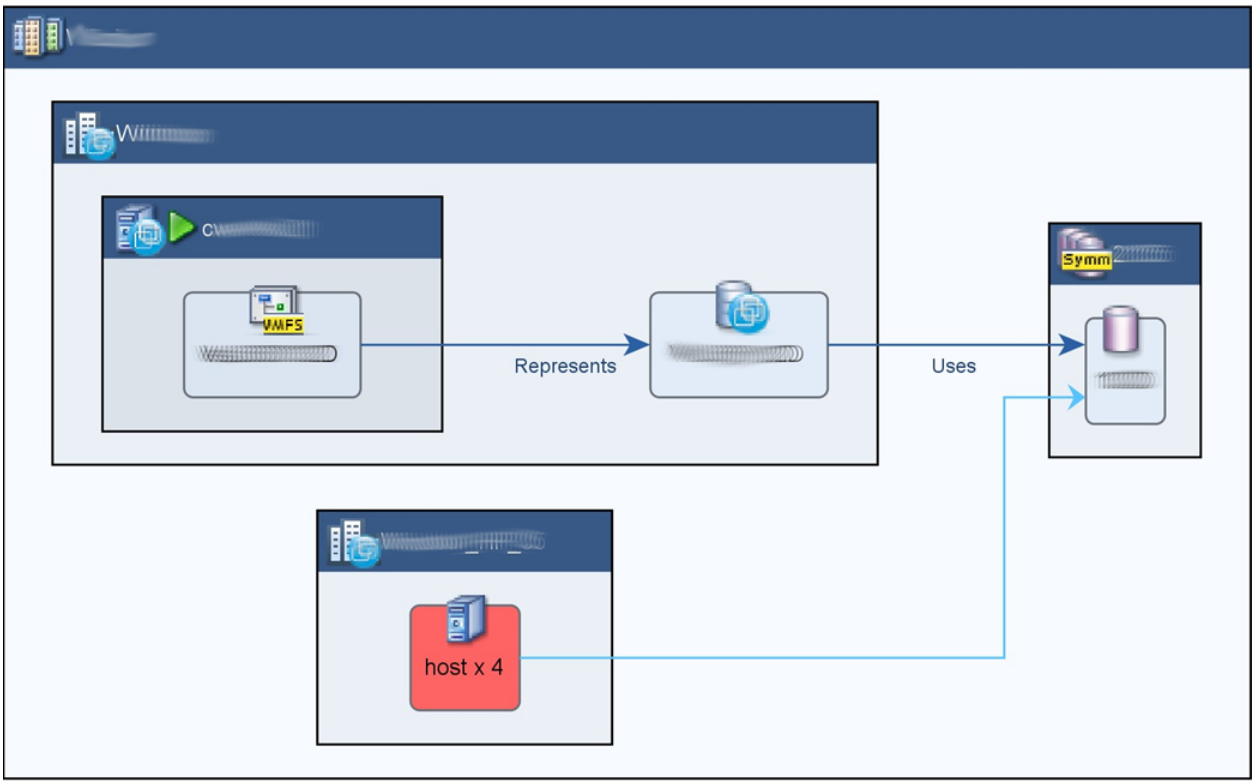

TICKET 46

Gap ID	00242PRIMARY	**Name**	Host accessing remote primary devices
Severity	WARNING	**Status**	OPEN
Categories	Optimization	**Rating**	
Detected on	Jan 23, 2012 1:11:53 PM		
Verified on	Jan 30, 2012 12:50:47 AM		

SUMMARY

ESX cluster node **VC-9b4** of ESX cluster **HP-BL-680** at site **Tokyo** is utilizing storage volumes located at site **Kofu**.

DESCRIPTION

ESX cluster node **VC-9b4** of ESX cluster **HP-BL-680** at site **Tokyo** is utilizing storage volumes located at site **Kofu**. This might result in performance degradation and unexpected outage.

The following table identifies the storage volumes in use by ESX cluster node VC-9b4 of ESX cluster HP-BL-680 at site Tokyo:

LV/VG/FS/DB	Physical volume	Storage volume	Storage array site
VMFS datastore **VMAX_S70**	**EMC Fibre Channel Disk (naa.6d37)**	**4001 /726**	**Kofu**
VMFS datastore **VMAX_S112**	**EMC Fibre Channel Disk (naa.6e46)**	**4001 /73B**	**Kofu**

IMPACT

The host may experience sub-optimal performance. Also, additional load on the network is generated, which potentially can be avoided and may impact the performance of other replication sessions.

Finally, unless the configuration is intentional, then any planned or unplanned outage of site London might lead to unexpected errors or failure of ESX cluster node **VC-9b4** of ESX cluster **HP-BL-680**.

RESOLUTION

Unless configured this way by design, migrate the storage volumes used to the local site. Pay attention to using the right storage tier and performance characteristics for the new storage volumes.

Finally, remember to check if scripts, scheduled tasks and array consistency group definitions are involved in the current configuration, as these will also require adjustment if used.

Note: If the host or array is not located at the site specified in the ticket, use the Configuration Wizard to correct the current site definition by associating the host or array with the correct site.

RESOLUTION

Similar tickets were opened for additional hosts.

TOPOLOGY

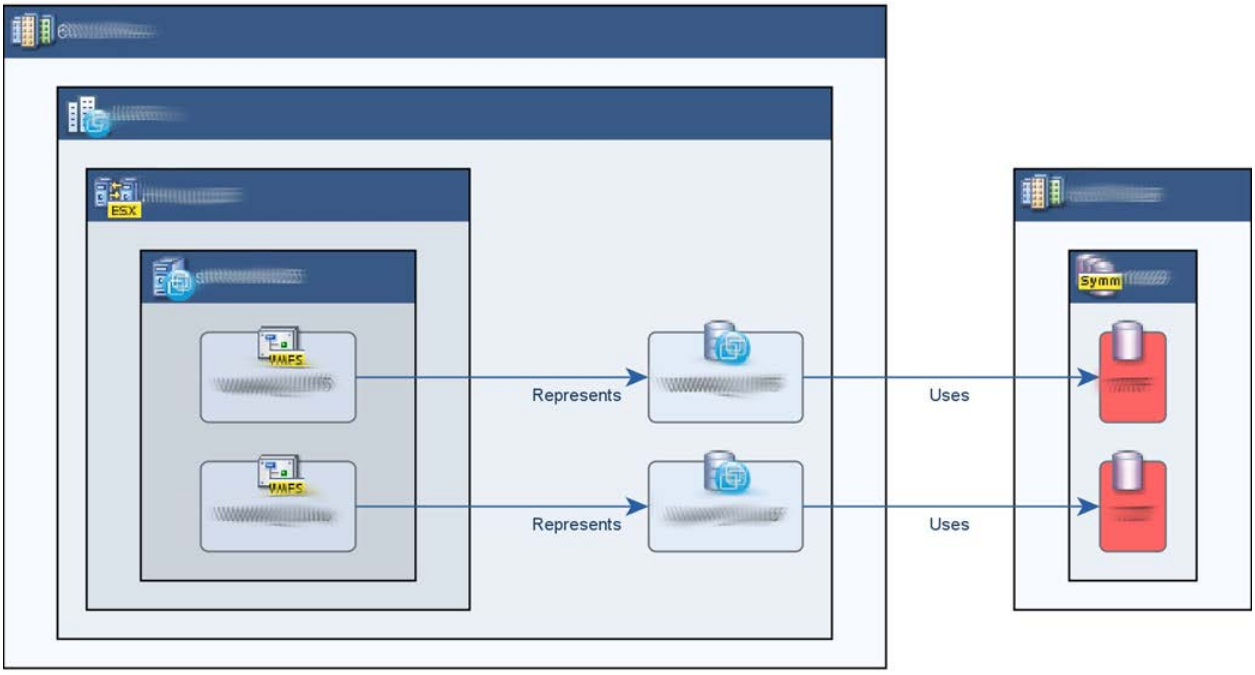

TICKET 28

Gap ID	00318SANIOSHSTG	**Name**	Hosts accessing shared storage with different number of SAN I/O paths
Severity	ERROR	**Status**	OPEN
Categories	Availability, Performance	**Rating**	
Detected on	Jan 23, 2012 12:10:23 PM		
Verified on	Jan 30, 2012 12:47:55 AM		

SUMMARY

The shared Symmetrix devices storing datastore **PCB6** of datacenter **LAB** at site **Tokyo** are accessed with different number of SAN I/O paths.

DESCRIPTION

The shared Symmetrix devices storing datastore **PCB6** of datacenter **LAB** are accessed by 2 nodes of the ESX cluster **HP-BL-680**. These cluster nodes are accessing these Symmetrix devices with different number of SAN I/O paths. This might result in performance degradation and increased risk of unexpected downtime (see impact).

Datastore PCB6 of datacenter LAB

Physical volume	Source Symmetrix device	Number of paths from VC-9b5	Number of paths from VC-9b4
EMC Fibre Channel Disk (naa.6d37) on 2 hosts	**4001 /727**	1	4

IMPACT

Having cluster nodes that access the same shared storage with different number of SAN I/O paths might result in performance degradation upon failover or switchover, since some cluster nodes have more paths than others. Since some of the cluster nodes are accessing the storage volumes with a single I/O path, it creates a single-point-of-failure, which might result in unexpected downtime and extended recovery time.

RESOLUTION

- Determine how many paths are required for the SAN storage volumes
- Add paths to the devices with inadequate configuration
- Strive to ensure end-to-end path redundancy (e.g., using separate array ports, fabric switches and host HBAs) as much as possible

TICKET 230

Gap ID	00490ESXPVLUNNUM	**Name**	ESX server physical volume has inconsistent LUN number
Severity	ERROR	**Status**	OPEN
Categories	Data Risk	**Rating**	★★★★★
Detected on	Nov 22, 2011 1:55:32 AM		
Verified on	Dec 28, 2011 1:50:10 AM		

SUMMARY

ESXi cluster node **prd-cl477** of ESX cluster **wss-cl-prd** at site **New York** has physical volume managed by VMware MPIO which has inconsistent LUN number.

DESCRIPTION

When ESX server uses SAN devices with VMware MPIO (with multiple HBA ports), it is important that the LUN number will be consistent between the physical paths, otherwise, it might lead to data corruption (see impact).

LUN number inconsistency was detected on ESXi cluster node **prd-cl477** of ESX cluster **wss-cl-prd** at site **New York**, which uses SAN devices with inconsistent LUN number.

The following table lists the physical volumes of prd-cl477 and the LUN number for each physical path:

Physical path	LUN number
Datastore **DMX-01** on Pseudo PV **EMC Fibre Channel Disk (sym.5643)** on Symmetrix device **4678 /0D4**	
fc.34f:1e4f-fc.559:559-sym.5643	24
fc.34f:1e4f-fc.556:556-sym.5643	23

The following VMs are affected by the gap:

- Windows VM **x1**
- Windows VM **x2**
- Windows VM **x3**
- Linux VM **l56**
- Linux VM **l57**
- Linux VM **l58**
- Linux VM **l309**
- Linux VM **l310**
- Linux VM **l311**
- Linux VM **l503**

IMPACT

When ESX server is using VMware MPIO with different HBA ports and the LUN number is inconsistent, there is a risk of data corruption. Since the ESX server is unaware that it uses the same SAN volume, it might treat it as different SAN volume and corrupt it's on data.

Resolution

Deactivate/Unmount the datastore/file system. Remap the LUN with the correct LUN number. Rescan for new devices. Activate the datastore.

TOPOLOGY

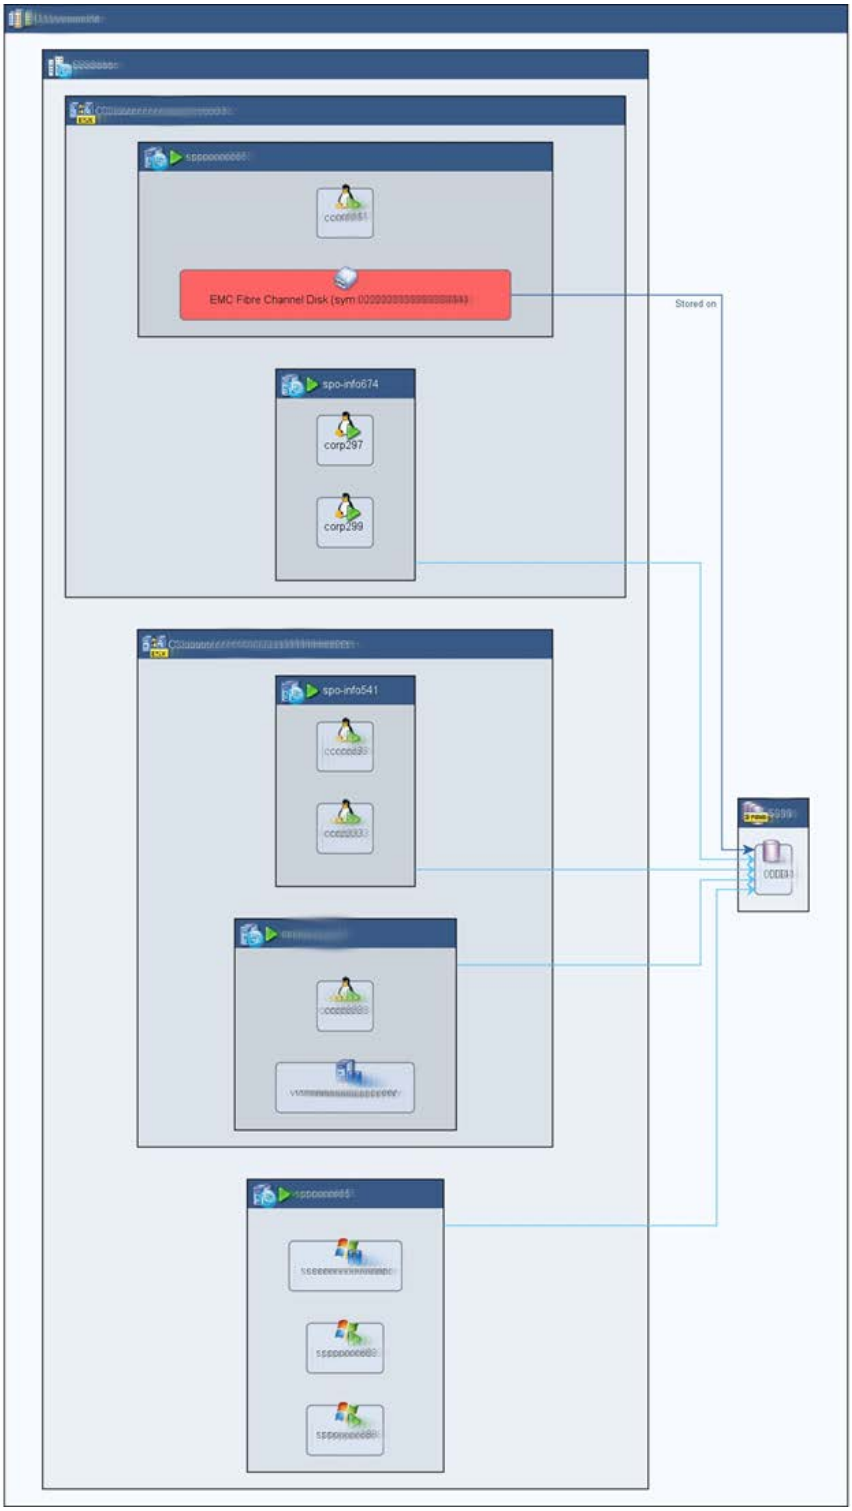

TICKET 92

Gap ID	00322SANIOSHSTG	**Name**	Hosts accessing shared storage with different SAN I/O policy
Severity	WARNING	**Status**	OPEN
Categories	Performance	**Rating**	★★★★☆
Detected on	Aug 04, 2011 8:22:24 AM		
Verified on	Sep 27, 2011 4:06:59 PM		

SUMMARY

The shared HDS logical units storing TrueCopy copy of datastore **RP-prd-sap3** of datacenter **Milan** at site **Milano** are accessed with different SAN I/O policy.

DESCRIPTION

The shared HDS logical units storing TrueCopy copy of datastore **RP-prd-sap3** of datacenter **Milan** are accessed by 2 nodes of the ESX cluster **PPRD**. These cluster nodes are accessing these HDS logical units with different SAN I/O policy. This might result in performance degradation (see impact).

TrueCopy copy of datastore RP-prd-sap3 of datacenter Milan:

Physical volume	Source HDS logical unit	Replica HDS logical unit	ravel82 I/O policy	ravel80 I/O policy
HITACHI Fibre Channel Disk (naa.60db5) on 2 hosts	**2306 /1238**	**3901 /1238**	VMW_PSP_RR (2 paths)	VMW_PSP_FIXED (2 paths)

Additional information to the table above:

- **Replication layout path:** TrueCopy

IMPACT

In addition to achieving redundancy, SAN I/O path policy may also be used to achieve higher I/O capacity through load balancing. The I/O policy is usually chosen to match the usage profile of the storage volumes. Having different I/O policy between cluster nodes might result in performance degradation upon failover or switchover, since at least one of the cluster nodes is not configured with the optimal configuration.

RESOLUTION

- Determine the I/O policy required for the SAN storage volumes
- Configure this policy to the devices with inadequate configuration
- Strive to ensure end-to-end path redundancy (e.g., using separate array ports, fabric switches and host HBAs) as much as possible

TICKET 43

Gap ID	00463VCTRDRSFA	**Name**	vCenter Management Server with fully automated DRS
Severity	INFO	**Status**	OPEN
Categories	Best Practice	**Rating**	★★★★★
Detected on	Aug 30, 2011 8:20:11 PM		
Verified on	Aug 31, 2011 7:07:39 PM		

SUMMARY

The VM **pd45-b-6** running vCenter Management Server is configured with fully automated DRS.

DESCRIPTION

If vCenter Management Server has been installed on a virtual machine within an ESX cluster it manages, it is recommended to disable DRS for that virtual machine, to prevent extended outage of vCenter upon certain failure scenarios (see impact).

The ESX cluster **MA-CL-602** at **New York** is managed by vCenter installed on the virtual machine **pd45-b-6**. This virtual machine is configured with fully automated DRS.

IMPACT

When the virtual machine running vCenter is down, it is impossible to manage the datacenters in its scope, Therefore, it is important to be able to start this virtual machine as fast as possible. If the virtual machine running vCenter is configured with fully automated DRS, it may unexpectedly migrate between cluster nodes, without the awareness of the VMware administrator. In this case, finding the virtual machine involves connecting to each ESX host in the cluster until the vCenter Management Server is found.

RESOLUTION

Disable DRS for the virtual machine running vCenter Management Server by changing the automation level.

In addition, it is recommended to:

- Document where the vCenter Server is located.
- Enable HA for the virtual machine running the vCenter Server, and setting the startup priority to "high".
- Make sure other services and servers on which vCenter depends are also starting automatically, with a high priority and in the correct order (Active Directory, DNS, SQL, others...).
- Make sure that the virtual machine running vCenter gets enough resources by setting the shares for both Memory and CPU to "high".

TICKET 185

Gap ID	00454DSLEN	**Name**	Disallow Snapshot LUN option disabled on ESX host
Severity	WARNING	**Status**	OPEN
Categories	Downtime, Best Practice, Data Risk	**Rating**	★★★★★
Detected on	May 25, 2012 4:23:45 PM		
Verified on	Jun 16, 2012 1:33:06 AM		

SUMMARY

The LVM.DisallowSnapshotLun option is disabled on ESX cluster node **PD-CL-54** of ESX cluster **ST-NT-04** at site **NEDC-12**.

DESCRIPTION

The **LVM.DisallowSnapshotLun** option controls whether an ESX server is allowed to mount storage volumes as VMFS even if the disk signature do not match. This option is usually used when recovering VMFS from Point-in-Time copies.

By default this option is enabled, to prevent duplicate allocation of the same VMFS from two (or more) different disks. A gap has been detected where this option is disabled on ESX cluster node **PD-CL-54** of ESX cluster **ST-NT-04** at site **NEDC-12**. This might lead to unexpected behavior in the ESX configuration (see impact).

Virtual machines affected by this gap:

- Windows VM **wmsv45**
- Linux VM **cig5l055**
- Windows VM **wmsv654**
- Windows VM **wmsv655**
- Windows VM **vt772**
- Windows VM **vt774**
- Linux VM **lndoraw2**
- Linux VM **lndora57**
- Linux VM **hj875**
- Linux VM **jko654**
- Linux VM **bdf46**
- Linux VM **ts4567**

IMPACT

Keeping **LVM.DisallowSnapshotLun** disabled (set to "0") on an ESX host might result in conflicting storage access and data loss when Point-in-Time copies are created at the storage layer. Since both source and target (copy) disks LUNs have the same disk ID, ESX will map each pair to the same VMFS. An attempt to write to the target will result in data corruption of the source.

LVM.DisallowSnapshotLun is usually changed from the default enabled state ("1") during maintenance, when access to storage based Point-in-Time copies is required for recovery purposes. It is important to reset it back to the "enabled" state as soon as possible.

RESOLUTION

Make sure that possible maintenance work has been completed and change the value of the option **LVM.DisallowSnapshotLUN** back to "1".

TOPOLOGY

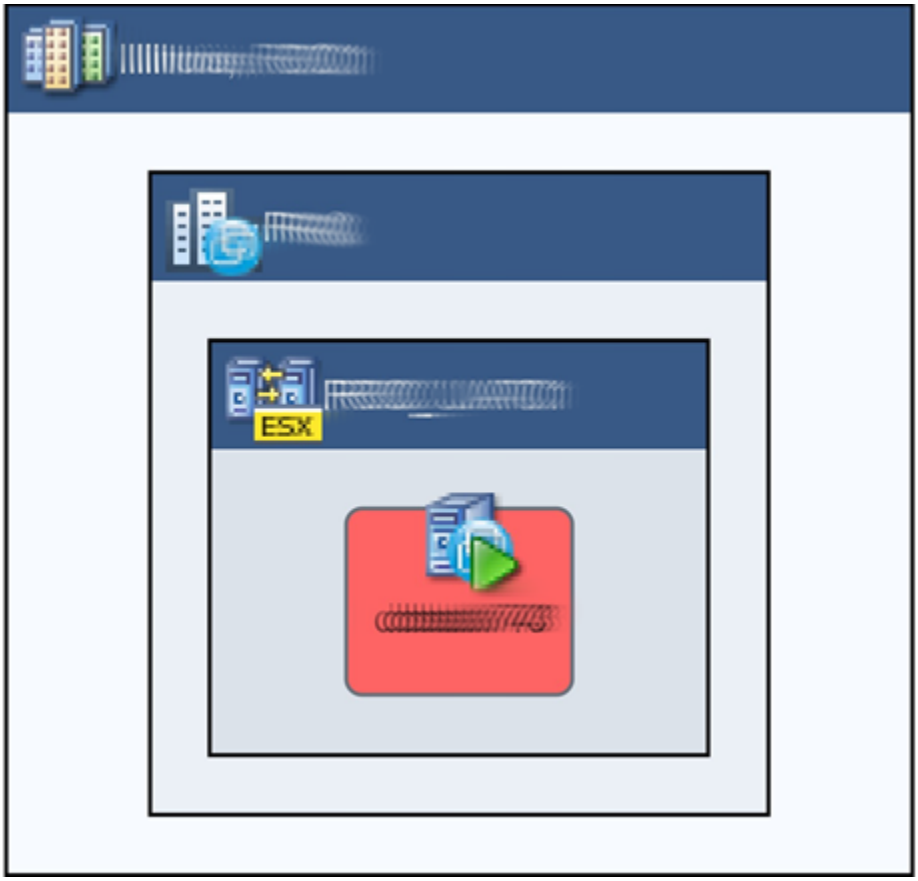

TICKET 42

Gap ID	00225DBF	**Name**	Mixture of database files
Severity	INFO	**Status**	OPEN
Categories	Best Practice	**Rating**	
Detected on	Jan 23, 2012 1:00:54 PM		
Verified on	Jan 30, 2012 12:40:52 AM		

SUMMARY

Local file system **C:** on Windows VM **rgfab** at site **Tokyo** contains a mixture of database files.

DESCRIPTION

Having different database file sets stored on the same storage volumes may result in incomplete copy of one or more of the databases, if future use of Point-in-Time copies is planned for these databases. In addition, performance issues may arise. Furthermore, it is recommended to store each database on a separated set of storage volumes.

The following table lists the database file-sets that are stored on local file system C:\ on Windows VM rgfab:

Database	File Type
Oracle database **ora_v_27** on Windows VM **rgfab**	Control files
Oracle database **ora_v_27** on Windows VM **rgfab**	Log files
Oracle database **ora_v_27** on Windows VM **rgfab**	Data files

The storage volumes that are storing the database files are not replicated.

IMPACT

If different file types are stored on the same storage volumes, they would compete over the same I/O resources, possibly resulting in suboptimal database performance. Moreover, since data files and transaction log files are the two most I/O intensive components of the database, it is a best practice to store different file types on separated storage volumes and spindles. This minimizes contention for the logs as new writes come in from the database and any old transaction log information is streamed out during incremental transaction log backups. It also isolates the sequential write and random read activity for these members from other volumes with differing access characteristics.

In addition, if both data and archived transaction log files are stored on the same storage volumes, then safe use of Point-in-Time copies is prevented, since:

- Some transaction logs will not be available for recovery
- The database copy might not recover successfully, since a log switch is required after data file replication ends and before archived transaction log replication starts. When both data files and archived transaction log files are stored on the same storage volumes, this best practice cannot be implemented. This might result in corrupt database blocks (sometime referred to as "split block") which may go unnoticed.

RESOLUTION

Consider separating your databases into different groups of storage volumes, thus eliminating data corruption risk and optimizing performance.

TOPOLOGY

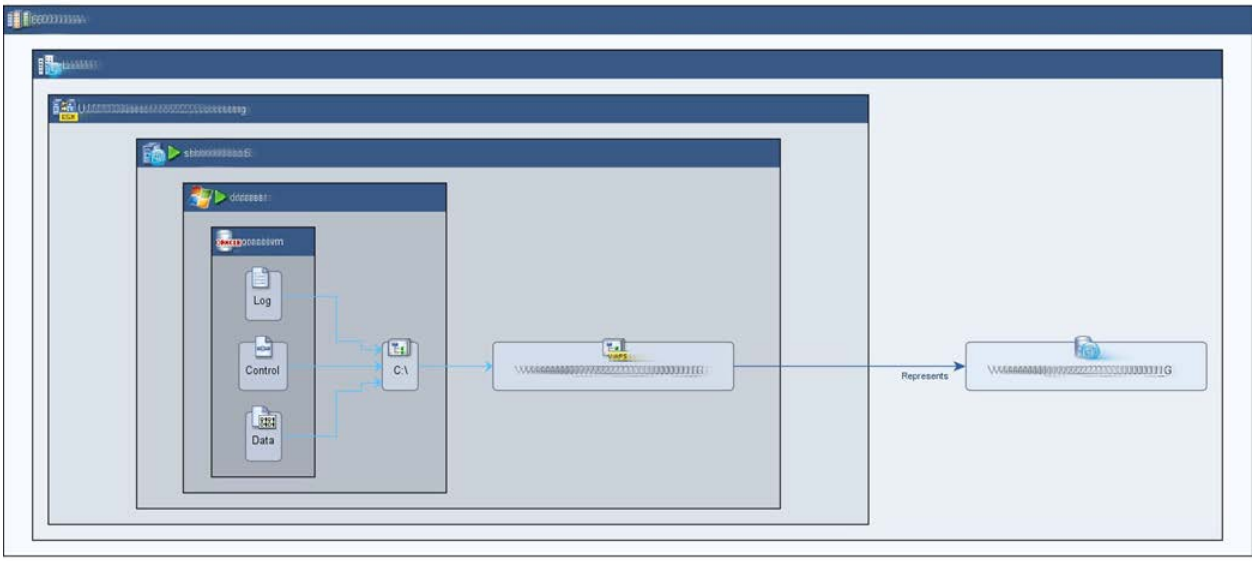

TICKET 173

Gap ID	00300STRGSPLTPS	**Name**	Storage volume physical split
Severity	INFO	**Status**	OPEN
Categories	Performance, Best Practice, Consistency	**Rating**	
Detected on	May 25, 2012 4:18:58 PM		
Verified on	Jun 16, 2012 1:30:55 AM		

SUMMARY

Storage array inconsistency was detected on Symmetrix devices storing VG **ws-prd-vg** on Linux VM **prd-ws-ne-56** at site **NEDC-12**.

DESCRIPTION

VG **ws-prd-vg** on Linux VM **prd-ws-ne-56** at site **NEDC-12** is stored on Symmetrix devices of different storage arrays.

This might result in degraded performance and data corruption or loss (see impact).

The following table lists the Symmetrix devices storing VG ws-prd-vg on Linux VM prd-ws-ne-56:

VM Physical volume/Disk	VMware Datastore	Source Symmetrix device	Storage array
Sd0 on **prd-ws-ne-56**	**PD-B6LUN02**	**2757 /1707**	**Symmetrix 2757**
Sd1 on **prd-ws-ne-56**	**PD-B7LUN37**	**4325 /124B**	Symmetrix **4325**
Sd2 on **prd-ws-ne-56**	**PD-B7LUN40**	**4325 /147C**	Symmetrix **4325**

IMPACT

Since the data is stored on several arrays, it will suffer downtime even when only one of the arrays is down.

If the data is replicated, upon split, the consistency of the replicas is not ensured for Symmetrix devices that are not a part of a consistency group (Composite group with enabled consistency mode).

Business impact:

- SAPR3

RESOLUTION

Review the mappings to the Symmetrix devices and consider remapping them on the same storage array.

Otherwise, consider using consistency group (Composite group with enabled consistency mode) on those storage volumes.

TOPOLOGY

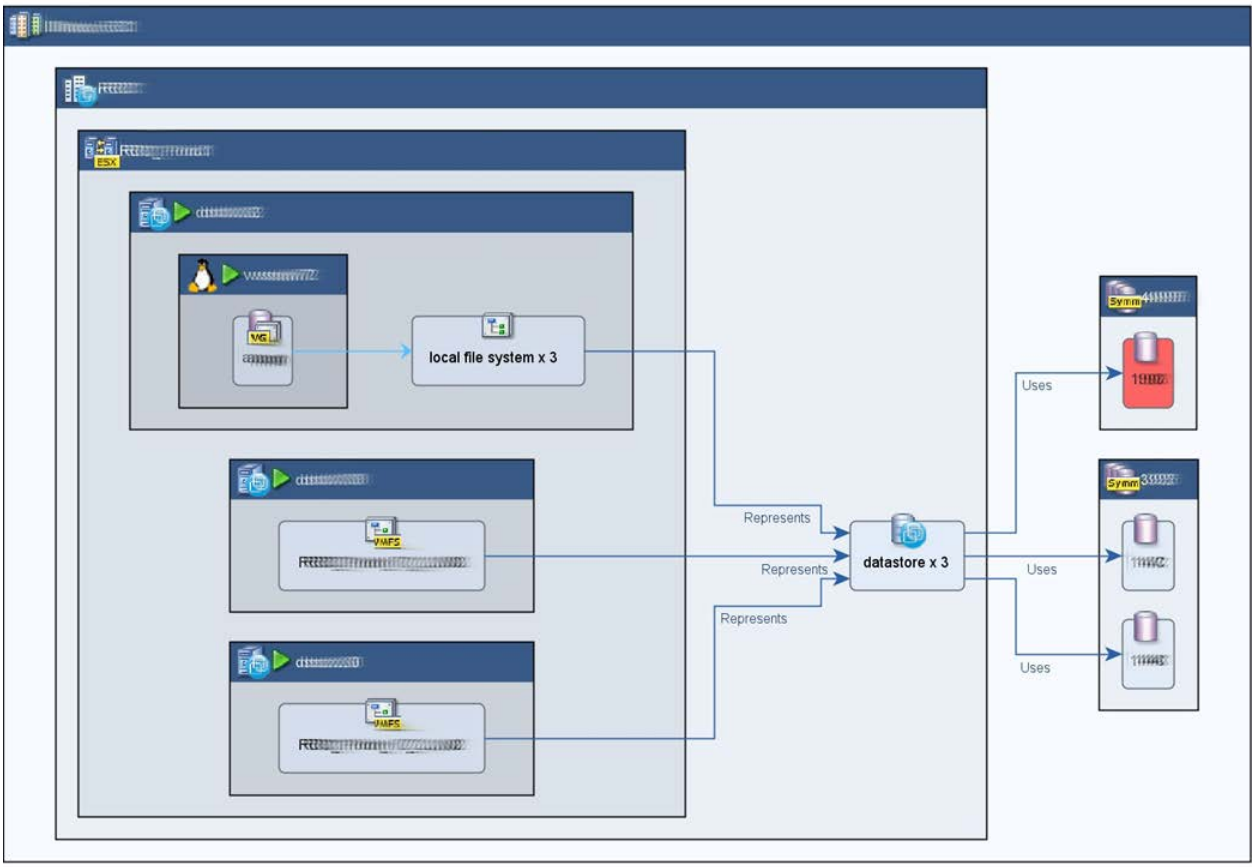
local file system x 3
Represents
Represents
Represents
datastore x 3
Uses
Uses
Uses

TICKET 211

Gap ID	00397HBASO	**Name**	Suboptimal HBA configuration
Severity	WARNING	**Status**	OPEN
Categories	Availability, Performance	**Rating**	★★★☆☆
Detected on	Nov 22, 2011 1:50:37 AM		
Verified on	Dec 28, 2011 1:52:10 AM		

SUMMARY

I/O redundancy has a single-point-of-failure in the HBA level on ESX cluster node **sdf56** of ESX cluster **CLNY8** at site **New York.**

DESCRIPTION

To ensure end-to-end path redundancy you must use separate array ports, fabric switches and host HBAs. A gap has been detected where physical volumes of ESX cluster node **sdf56** of ESX cluster **CLNY8** at site **New York** are configured with multi-pathing, but they have a single point of failure in the HBA level.

This means that either the physical paths are mapped to a single HBA port or that multiple ports are in use, but these ports are on the same HBA device. The following section provides detailed information about these physical volumes.

The following physical volumes are storing datastore DMX-05 of datacenter New York at site New York:

Host	Pseudo PV	Symmetrix device	HBA port
ESX cluster node **sdf56** of ESX cluster **CLNY8**	**EMC Fibre Channel Disk (sym.5674)**	**4457 /421**	**vmhba1**

The following list displays the available HBAs on **sdf56**:

- **vmhba0**
- **vmhba1**
- **vmhba2**
- **vmhba3**

IMPACT

Having multiple paths using a single HBA device eliminates the multipathing in the host level. The data will not be available upon failure of a single HBA.

RESOLUTION

Make sure that each physical path (of each pseudo PV) uses a different HBA port device. Make sure that the HBA distribution is identical between volumes storing the same data (e.g. volume group). If necessary, install a new HBA device on the host.

TICKET 118

Gap ID	00260DBLFRD	**Name**	Database log files on suboptimal storage
Severity	INFO	**Status**	OPEN
Categories	Performance, Best Practice	**Rating**	
Detected on	Nov 22, 2011 1:57:37 AM		
Verified on	Dec 28, 2011 1:59:10 AM		

SUMMARY

Symmetrix devices storing log file of 3 databases (model, msdb, master) on Windows cluster node **spo-info396** of MSCS **MSCLUSTER05** are stored on suboptimal storage.

DESCRIPTION

To achieve optimal performance, it is a best practice to store database log files on RAID1, RAID0+1 or RAID1+0. The RAID level of the Symmetrix devices storing the log files of 3 (model, msdb, master) on Windows cluster node **spo-info396** of MSCS **MSCLUSTER05** do not meet this standard. This might lead to performance degradation (see impact).

The following databases store their log files on these Symmetrix devices:

- SQL Server database **sql06 /master** on Windows cluster node **spo-info396** of MSCS **MSCLUSTER05**
- SQL Server database **sql06 /model** on Windows cluster node **spo-info396** of MSCS **MSCLUSTER05**
- SQL Server database **sql06 /msdb** on Windows cluster node **spo-info396** of MSCS **MSCLUSTER05**

The following table lists the Symmetrix devices and their RAID levels:

Physical volume	Source Symmetrix device	RAID level
Disk5 on **spo-info396**	**4678 /33F**	RAID5

IMPACT

Since the database log files are I/O intensive (both read and write), database vendors recommend RAID mirroring (for data protection) and striping (for performance). The recommended RAID levels are:

- RAID 1+0
- RAID 0+1
- RAID 1 (while not striped, it still offers significantly better performance than RAID5 and similar shared parity configurations)

Using other RAID levels might significantly reduce database performance.

RESOLUTION

Allocate storage volumes with a proper RAID level that provides high performance and configure the databases to store the log files on these devices.

TOPOLOGY

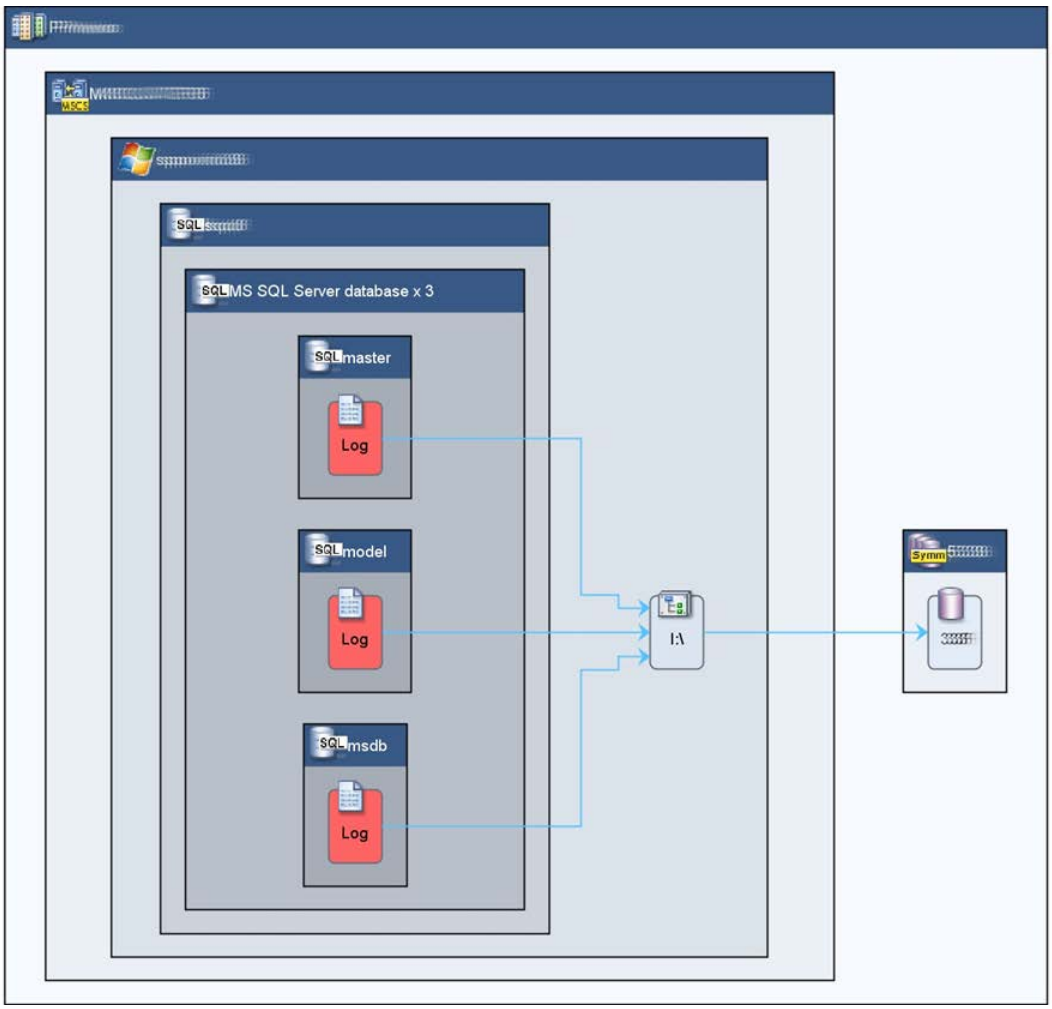

TICKET 36

Gap ID	00239PSWP	**Name**	Swap device on sub-optimal storage
Severity	INFO	**Status**	OPEN
Categories	Performance	**Rating**	
Detected on	Nov 22, 2011 1:59:37 AM		
Verified on	Dec 28, 2011 1:55:10 AM		

SUMMARY

Windows VM **ipfg5** at site **LA** has swap files stored on RAID6 devices.

DESCRIPTION

Windows VM **ipfg5** at site **LA** has swap files stored on Hitachi RAID6 devices. This configuration might seriously affect its performance (see impact).

The following swap files are defined on host ipfg5:

Swap File	VG	Size	Raid Level
C:\pagefile.sys	N/A	2GB	RAID6

Swap file C:\pagefile.sys is stored on the following RAID6 devices:

PV	Source Storage Volume	Raid Level
HITACHI Fibre Channel Disk (naa.668)	**2270 /0068**	RAID6

IMPACT

Swap files are extremely volatile and will typically involve huge amounts of I/O. For this reason, vendors will always suggest to store them on devices which offer the best possible I/O characteristics in terms of access time, I/O operations per second and throughput.

RAID5 devices offer significantly slower write performance than other storage options, such as unprotected (RAID0) or mirrored (RAID1) configurations. This might result in dramatical performance degradation of the host.

Business impact:

- ESX LA
- Exchange

RESOLUTION

Consider migrating the swap files from the current storage devices to unprotected or RAID1 devices.

TOPOLOGY

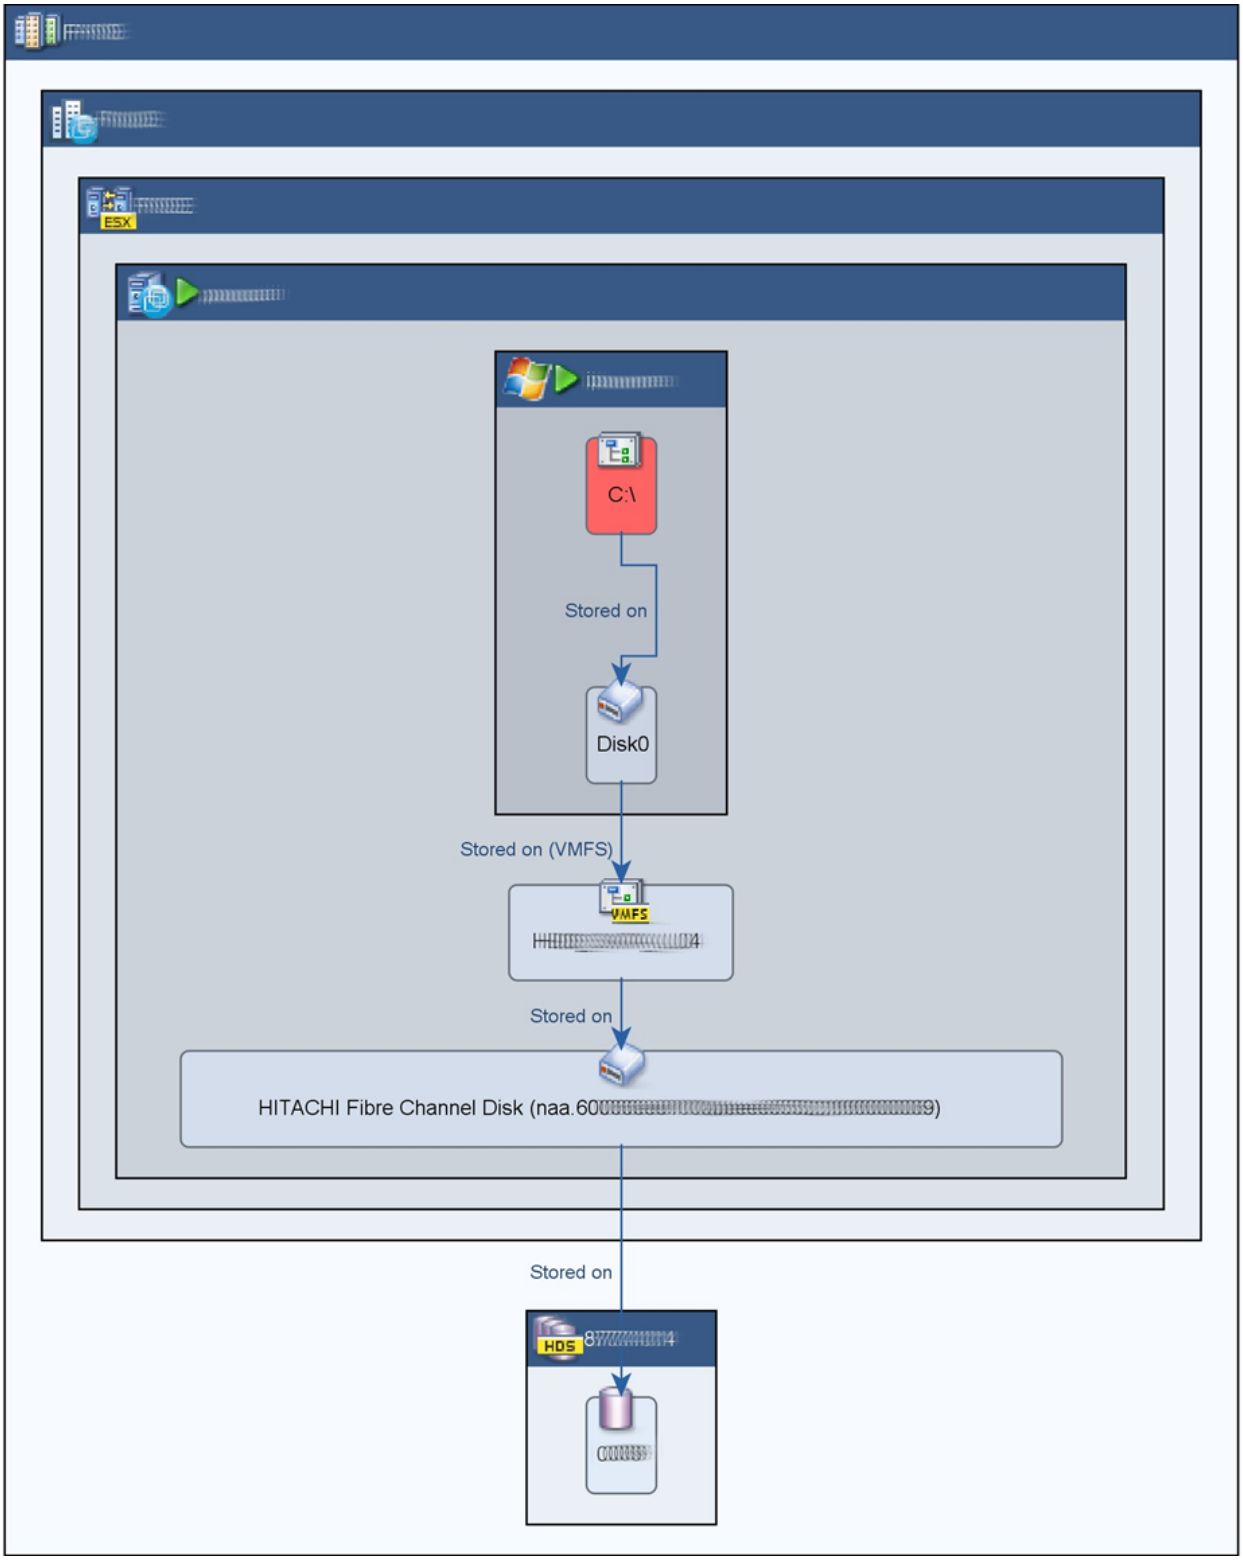
ESX
C:\
Stored on
Disk0
Stored on (VMFS)
VMFS
Stored on
HITACHI Fibre Channel Disk (naa.60
Stored on
HDS

NODE COMPARISON REPORT

OPERATING SYSTEM

Component	esx-nt-ew-101	esx-nt-ew-102	esx-nt-ew-103	esx-nt-ew-104
Version	4.0.0	4.1.0	4.1.0	4.1.0
Version Edition	**VMware ESXi 4.0.0 build-261974**	VMware ESXi 4.1.0 build-433742	VMware ESXi 4.1.0 build-433742	VMware ESXi 4.1.0 build-433742
Version Level	**261974**	433742	433742	433742

OPERATING SYSTEM OPTIONS

Option	esx-nt-ew-101	esx-nt-ew-102	esx-nt-ew-103	esx-nt-ew-104
FT.TCPNoDelayPrimary	**0**	1	1	1
FT.Vmknic	**vmk1**	vmk2	vmk2	vmk2
Mem.VMXMinMB	**32**	56	56	56
Misc.HostAgentUpdateLevel	**2**	1	1	1
Misc.TimerMaxHardPeriod	**2000**	100000	100000	100000
Net.MaxPageInQueueLen	**500**	75	75	75
Net.TcpipDefLROMaxLength	**16000**	65535	65535	65535
Syslog.Remote.Hostname	**N/A**	nyctsyslog	**N/A**	nyctsyslog