DR ASSURANCE MONTHLY REPORT – XXXX BANK – DECEMBER 2011

DR Assurance provided by

A total of 32 potential DR and HA vulnerabilities were detected by RecoverGuard's daily monitoring of the bank's multiple data centers covered by the DR Assurance service.

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

A total of 14 high priority threats identified by RecoverGuard are currently open and require attention. This is down from 15 issues that were open last month.

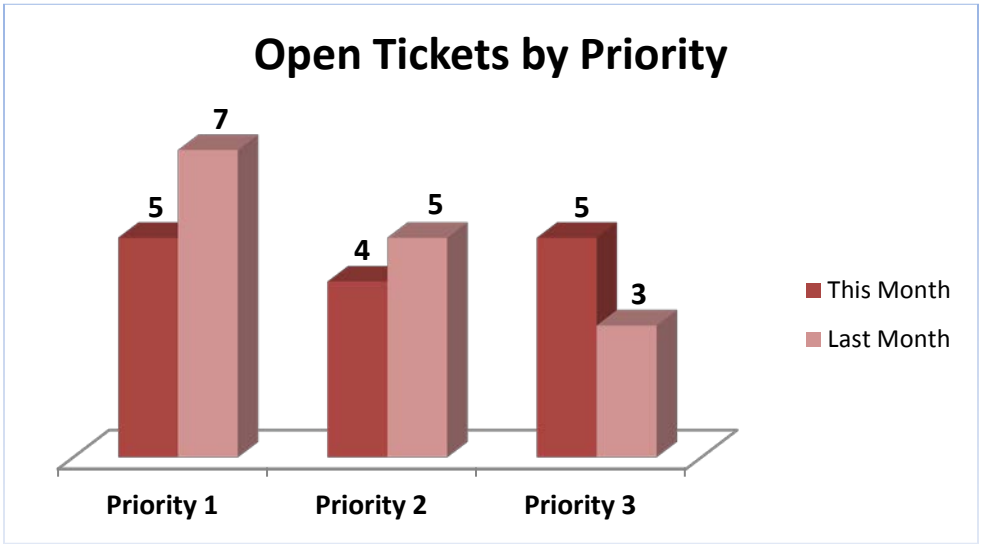

The chart below shows priority 1-3 issues by the business service impacted.

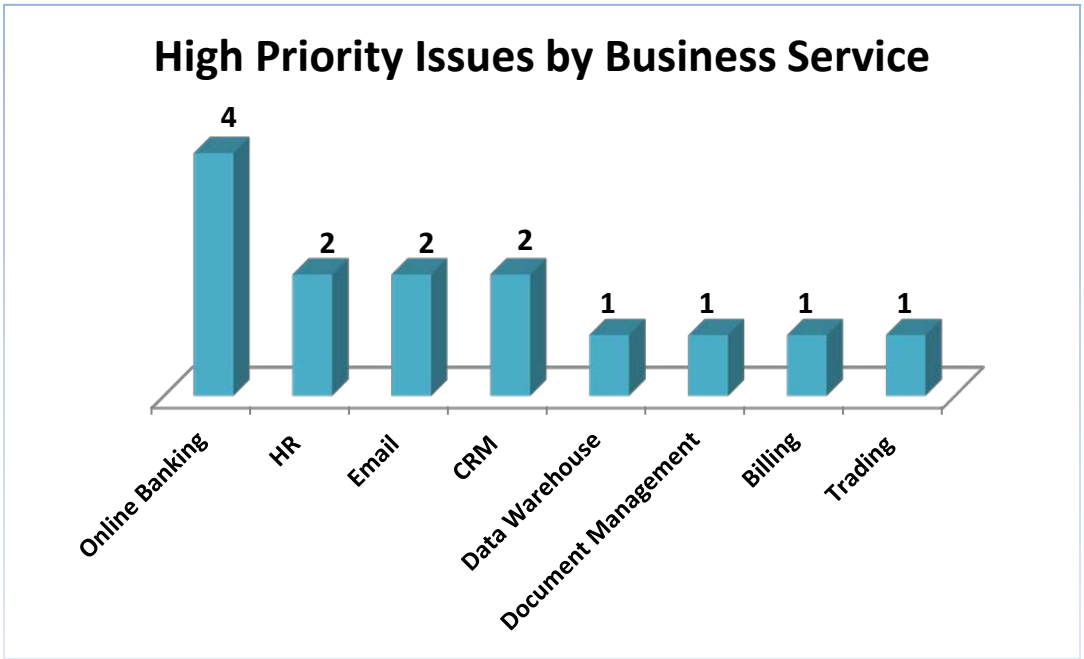

Most high priority open issues are related to replications.

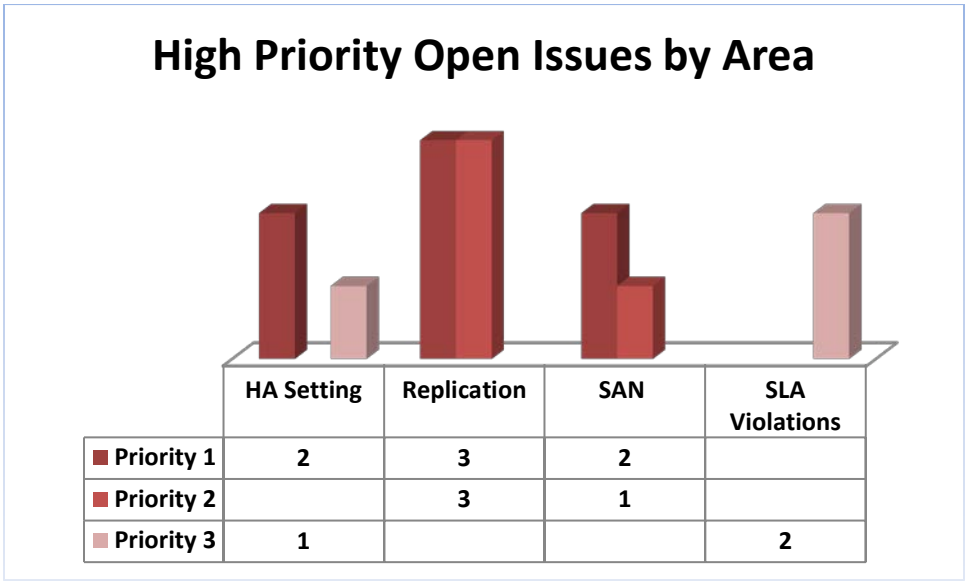

	HA Setting	Replication	SAN	SLA Violations
Priority 1	2	3	2	
Priority 2		3	1	
Priority 3	1			2

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 HA) • Potential data loss upon fail-over	1	2	1, 4
Online Banking	Replication	Partial replication	• Data loss in the event of a site outage • Downtime	1	2	2, 14
Trading	HA Setting	Inconsistent I/O capacity across nodes	• Unexpected performance degradation as cluster service fails over (or re-assigned) to a standby node	3	1	18
Billing	Replication	Replicas unmapped to DR servers	• Extended downtime in the event of site outage • Potential data loss	1	1	9
CRM	Replication	Inconsistent database backup	• Extensive data loss in the event of a physical or logical fault in the primary copy	2	1	7
CRM	SAN	Unauthorized storage access	• Risk of massive data corruption affecting multiple Virtual Machines	2	1	17
Email	SAN	Incorrect redundancy configuration	• single point of failure in a system designed for redundancy	1	1	20
Email	SAN	Insufficient number of hot-spare drives	• Insufficient redundancy might result in massive data loss affecting multiple servers	1	1	23
HR	SLA Violations	HA/DR Service Level Agreement breach	• No compliance with data retention requirement • Insufficient performance and redundancy expected after fail-over	3	2	5, 8
Document Management	Replication	Incorrect storage consistency group configuration	• Data loss in the event of a site outage • Downtime	2	1	10
Data Warehouse	Replication	Excessive Oracle DataGuard heartbeat failures	• RPO violation	2	1	27

DR TESTING AS A SERVICE – HOW DOES IT WORK?

The proactive continuous monitoring approach used by the DR & HA Assurance Service eliminates the inefficient manual testing and DR audits that consume IT resources, yet still fail to uncover many critical risks, data protection breaches and recoverability gaps.

Your DR & HA Assurance Service begins with the seamless integration of RecoverGuard into your IT infrastructure. The software operates in a **non-intrusive, read-only mode**. Within the first day of deployment:

- RecoverGuard begins scanning your storage, databases, servers and replication configurations for risks.
- Automatic analysis tools and our powerful Gap Detection Engine identify vulnerabilities such as unprotected databases or database partitions, noncompliant replication configurations, data that cannot be recovered to a valid consistency point, and much more.
- Daily statistics are sent to Continuity Software's DR experts for review and analysis. No business or confidential data is ever touched or sent outside of your organization.
- When a problem is detected, our experts will alert you with details regarding the severity of the problem and how to resolve it.
- Weekly, monthly and quarterly summary reports detail findings on issues uncovered and suggested resolutions, making them useful for both Business Continuity and Disaster Recovery professionals.

For more information, see http://www.continuitysoftware.com/services/drassurance.

Limited Time Offer: 30-day free trial - http://www.continuitysoftware.com/dr-service-trial

TICKET 1

Gap ID	00502VCSVISMIS	**Name**	VCS passive system not connected to devices
Severity	ERROR	**Status**	OPEN
Categories	Availability		
Detected on	Dec 04, 2011 12:41:29 AM		
Verified on	Dec 10, 2011 1:07:30 AM		

SUMMARY

Passive node of VCS ■■■■ at site ■■■■ is not connected to all required storage volumes

DESCRIPTION

A resource ■■■■ of service group ■■■■ in VCS ■■■■ refers to VG ■■■■, yet not all the passive nodes have SAN connectivity to the VG storage volumes. This might result in unexpected downtime upon failover or switchover (see impact).

The following table presents the missing host-to-storage volume connections:

Passive node	Storage volume
■■■■	IBM. ■■■■ /10A0

IMPACT

When a host does not have access to a required SAN volume, then the import of the VG that is stored on these storage volumes is doomed to fail. Upon failover or switchover the entire service group will fail, resulting in unexpected downtime.

Business impact: Online Banking

RESOLUTION

Reconfigure your SAN to allow access to the missing storage volumes by the required hosts. Remember to consider required performance and redundancy levels, as you may need to set more than one path to each missing storage volume. Reference your active nodes to determine how many paths are required and what multi-pathing mode to use.

TOPOLOGY

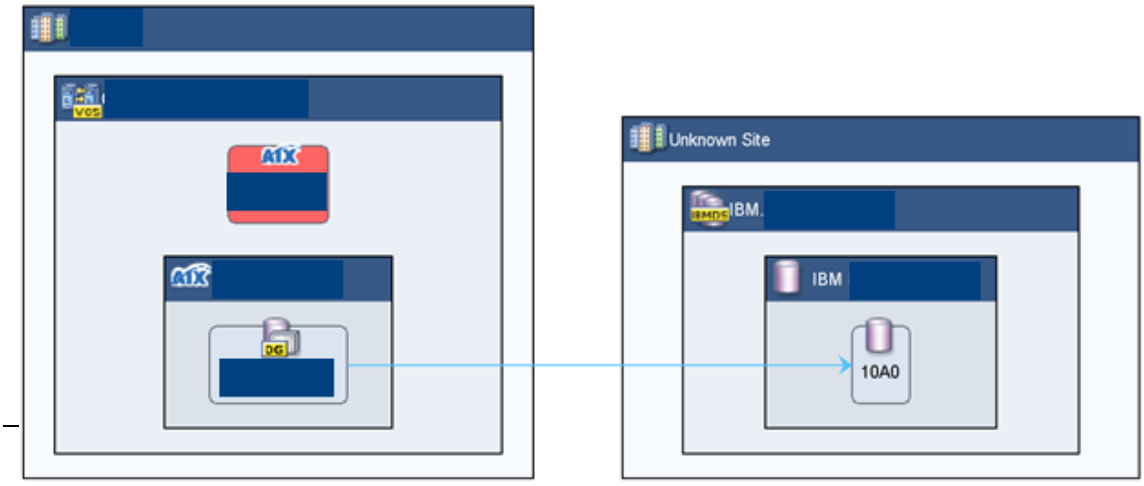

TICKET 2

Gap ID	00301RDTCONST	**Name**	Replica data inconsistency
Severity	ERROR	**Status**	OPEN
Categories	Extended Recovery Time, Consistency		
Detected on	Dec 20, 2011 5:05:32 AM		
Verified on	Dec 26, 2011 6:28:36 PM		

SUMMARY

Replication state inconsistency was detected on replicas of HDS logical units storing VG ■■■■ on AIX host ■■■■ at site ■■■■

DESCRIPTION

VG ■■■■ on AIX host **■■■■** is stored on HDS logical units that have the following replica sets:

- TrueCopy copy
- 2 TrueCopy-QuickShadow copies

Typically, all replicas in each replica set should be consistent in terms of replication state, last action and timings, to maintain a consistent and reliable copy. Lack of uniformity often suggests that data will be corrupt of lost upon disaster (see impact).

Data consistency problems were detected in 2 replica sets. The following tables provide more information on each set.

Replica set 1: TrueCopy-QuickShadow copy of VG ■■■■ on AIX host ■■■■ at site ■■■■

Physical volume	Source HDS logical unit	Replica HDS logical unit	Replication state
hdisk188	**■■■■ /1034**	**■■■■ /16BA**	**Pair**
hdisk179	**■■■■ /13E0**	**■■■■ /16BF**	**Split**
hdisk180	**■■■■ /1488**	**■■■■ /16BE**	**Split**

Additional information to the table above:

- **Replication layout path:** TrueCopy-QuickShadow
- **Replica devices are located at site: ■■■■**

Databases affected by this gap

- Data files of Oracle database ■■■■ on AIX host ■■■■

IMPACT

Typically, all replicas in each replica set should be consistent in terms of replication state, last action and timings, to maintain a consistent and reliable copy. Lack of uniformity often suggest that data will be corrupt or lost upon disaster since it will usually imply that some of the devices contain data which is more current in time than others. This usually means that data integrity is lost.

Business impact: Online Banking

RESOLUTION

Review existing mechanisms, such as scripts that are being used to manage the replica sets, to ensure all storage volumes in a set are split simultaneously.
Consider using device groups or composite groups to ensure consistency on split.
If some replicas entered an error state or were manually changed, correct the replication to achieve a valid replica set.

TOPOLOGY

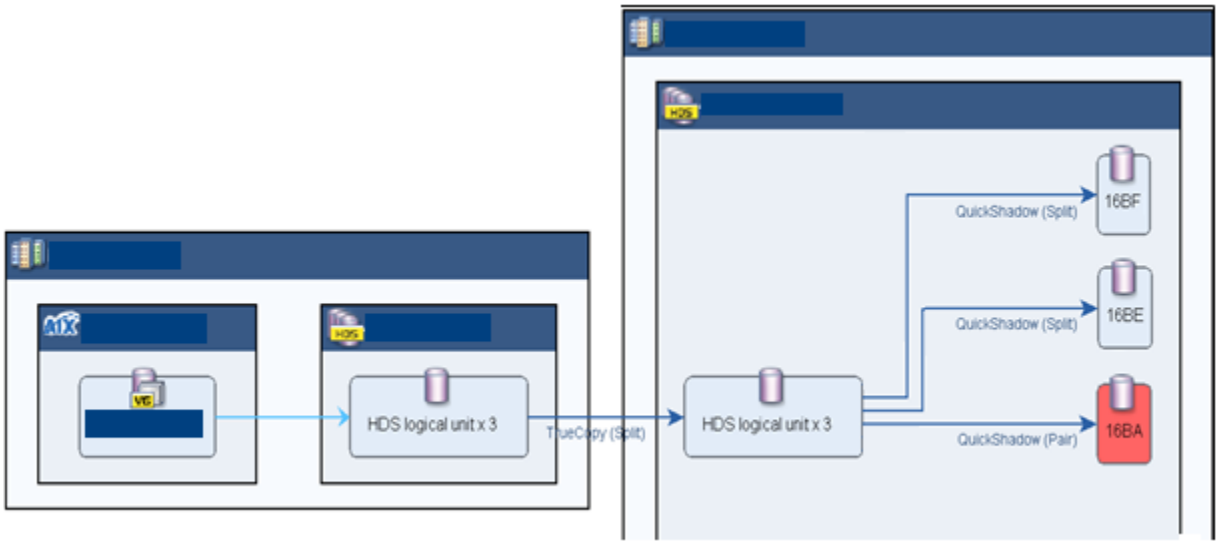
HDS logical unit x 3
TrueCopy (Split)
HDS logical unit x 3
QuickShadow (Split)
16BF
QuickShadow (Split)
16BE
QuickShadow (Pair)
16BA

TICKET 4

Gap ID	00500VCSOFFMNTFAIL	**Name**	VCS mount resource properties defined on offline system
Severity	ERROR	**Status**	OPEN
Categories	Availability		
Detected on	Dec 27, 2011 4:04:40 PM		
Verified on	Dec 27, 2011 4:04:40 PM		

SUMMARY

A file system has an incorrect mount configuration on a passive node of VCS ■■■■ at sites ■■■■ and ■■■■

DESCRIPTION

Cluster mount resource file systems configured on passive nodes must have the same mount directory and block device as those configured for the mount resource. Great care must be taken so that no local file system will accidentally be configured on a resource-related mount directory or block device.
A passive node of VCS ■■■■ does not meet these requirements.

The following hosts of VCS ■■■■ have a configuration conflict between local file systems and cluster mount resources:

Host	Service Group	Resource	Resource Block Device	System Block Device
■■■■	■■■■	■■■■_Mount	/dev/vx/dsk/dg_■■■■/vol_**prep**_oracle	/dev/vx/dsk/dg_■■■■/vol_**prod**_oracle

IMPACT

If the passive node should be able to take ownership of the mount resource upon failover, and its file system configuration does not have exactly the same mount point and block device specified for the mount resource, then failover would not succeed, resulting in unexpected downtime.
All other local file systems should not accidentally have either the mount point or block device identical to one of the cluster resources, since this might result in unexpected service disruption or data corruption.

Business impact: Online Banking

RESOLUTION

Review the mis-configured file systems and cluster mount resources and perform the following;

- For cluster mount resources - make sure that the file system mount point and block device match those defined for the resource
- For all other file systems - correct the mis-configured mount point or block device, or consider removing the file system if not required

TICKET 5

Gap ID	00701NTPSLARET	**Name**	NetApp filer violating retention SLA
Severity	ERROR	**Status**	OPEN
Categories	Data Risk, Replication SLA		
Detected on	Dec 31, 2011 2:08:42 PM		
Verified on	Dec 31, 2011 7:06:23 PM		

SUMMARY

The NetApp filer ■■■■ violates the retention SLA defined in the policy Snapshots SLA

DESCRIPTION

The NetApp filer ■■■■ at ■■■■ was assigned the SLA policy **Snapshots SLA**. When assigning a NetApp filer, volume or Q-Tree to SLA policy, the assigned component must meet all the criteria configured in the SLA policy. However, SLA violation was detected in the context of **retention.**

The following table shows the exciting snapshots on each volume, emphasizing that there are SLA breaches in the amount of snapshots on the following volumes:

Affected component	Violated SLA definition	De-facto
NetApp volume /vol/■■■■	**40 local logically protected copies with daily frequency, replication path: Snapshot**	* Snapshot (age: 02:07:42 hours) * Snapshot (age: 08:07:42 hours) * Snapshot (age: 18:07:41 hours) * Snapshot (age: 22:07:42 hours) * Snapshot (age: 1 day and 02:07:42 hours) * Snapshot (age: 1 day and 08:07:42 hours) * Snapshot (age: 1 day and 18:07:41 hours) * Snapshot (age: 1 day and 22:07:42 hours) * Snapshot (age: 2 days and 18:07:41 hours) * Snapshot (age: 3 days and 18:07:41 hours) * Snapshot (age: 4 days and 18:07:41 hours)
NetApp volume /vol/■■■■	**40 local logically protected copies with daily frequency, replication path: Snapshot**	* Snapshot (age: 02:07:42 hours) * Snapshot (age: 06:07:42 hours) * Snapshot (age: 10:07:42 hours) * Snapshot (age: 18:06:41 hours) * Snapshot (age: 22:07:42 hours) * Snapshot (age: 1 day and 02:07:42 hours) * Snapshot (age: 1 day and 06:07:42 hours) * Snapshot (age: 1 day and 18:06:41 hours) * Snapshot (age: 2 days and 18:06:41 hours) * Snapshot (age: 3 days and 18:07:41 hours) * Snapshot (age: 4 days and 18:07:41 hours)

IMPACT

Failing to meet the required retention means that in certain failure scenarios, no valid point in time copies will be available for fast recovery (for example, a logical data corruption might propagate to all existing copies before identified, since not enough history is retained).

Business impact: HR

RESOLUTION

If the SLA deviations are by design, the resolution can be one of the following:

- Re-configure the SLA policy to be more accurate regarding the desired SLA
- Assign the SLA policy to specific volumes or Q-Trees
- Suppress the SLA deviations that are by design, using the **Symptoms** tab in the Tickets module

If the SLA deviations are not by design, re-configure the filer or replicas so they will meet the SLA policy defined to them. Use the following section as a reference to fix the SLA issues in your environment:

The SLA policy Snapshots SLA is defined as followed:

Data Protection definition:

- RPO was not defined -

Retention:

- 40 local logically protected copies with daily frequency, replication path: Snapshot (**Enforced**)

TICKET 7

Gap ID	00241IDIHC	**Name**	Inconsistent database image - Hot backup and storage consistency solution missing
Severity	WARNING	**Status**	OPEN
Categories	Data Risk, Best Practice		
Detected on	Dec 15, 2010 6:48:26 AM		
Verified on	Dec 23, 2011 4:55:23 AM		

SUMMARY

SRDF/Clone replicas of Oracle database ■■■■ on HPUX cluster node ■■■■ of cluster ■■■■ at site ■■■■ are not activated/split using database or storage-based consistency mechanisms

DESCRIPTION

Host ■■■■ is storing Oracle database ■■■■ data files on Symmetrix devices which have SRDF/Clone replicas that were activated/split:

- While the database was online
- But not during hot backup mode
- And with no Symmetrix consistency assurance

This might result in data loss upon disaster or data restore.

The following table lists the source and replica SRDF/Clone devices of Oracle database ■■■■ on host ■■■:

VG/LV	PV	Source devices	Target devices	Layout path	Data Point-in-Time	In backup mode
■■■■	**c16t1d2**	■■■■/2B8	■■■■/86C	Clone,SRDF/S,Clone	Feb 27, 2011 10:57:56 PM	No
■■■■	**c14t2d4**	■■■■/2E0	■■■■/894	Clone,SRDF/S,Clone	Feb 27, 2011 10:57:56 PM	No
■■■■	**c16t2d3**	■■■■/2DC	■■■■/890	Clone,SRDF/S,Clone	Feb 27, 2011 10:57:56 PM	No

34 rows were filtered

Additional information: Notice that the hot back up mode is not synchronized with the replication process at any time as it should be. This might cause an inconsistency in the database backup.

- Last hot backup of Oracle tablespaces: ■■■■, ■■■■, ■■■■, ■■■■, ■■■■, ■■■■, ■■■■:
  Began at: Feb 27, 2011 9:18:50 PM
  Ended at: Feb 27, 2011 9:19:20 PM
- Last hot backup of Oracle tablespaces: ■■■■, ■■■■, ■■■■, ■■■■:
  Began at: Feb 26, 2011 9:18:30 PM
  Ended at: Feb 26, 2011 9:19:20 PM
- Last hot backup of Oracle tablespaces: ■■■■, ■■■■, ■■■■, ■■■■, ■■■■, ■■■■:
  Began at: Feb 26, 2011 9:16:30 PM
  Ended at: Feb 26, 2011 9:19:20 PM

IMPACT

The consistency of the database image is not guaranteed. If transactions were committed while copy or split was performed, it is possible that the database image is inconsistent, thus cannot be used for recovery, and data might be lost.

Business impact: CRM

RESOLUTION

Synchronize the timing of Oracle hot backup and the timing of replica.

TICKET 8

Gap ID	00700HSLASANIO	Name	Host violating minimum of SAN I/O paths SLA
Severity	ERROR	Status	OPEN
Categories	SLA Breach, Performance, Availability		
Detected on	Dec 05, 2011 7:49:58 AM		
Verified on	Dec 13, 2011 6:07:24 PM		

SUMMARY

The AIX host ■■■■ violates minimum number of I/O paths SLA defined in the policy BC Tier 6 - Silver

DESCRIPTION

The AIX host ■■■■ at site ■■■■ was assigned the SLA policy "BC Tier 6 - Silver ". When assigning a host, VG, file system or database service to a SLA policy, the assigned component must meet all the criteria configured in the SLA policy. However, an SLA breach was detected in the context of minimum number of I/O paths.

The following table specifies the minimum number of I/O paths SLA definition and the SLA de-facto:

Affected component	Violated SLA definition	De-facto
VG ■■■■	Minimum 2 I/O paths	1 I/O path
VG ■■■■	Minimum 2 I/O paths	1 I/O path

The following table presents the I/O paths information:

Source VG	Physical volume	Source Symmetrix device	SAN paths information
■■■■	hdiskpower3	■■■■/1B8	hdisk12 (alive)
■■■■	hdiskpower2	■■■■/1B7	hdisk11 (alive)
■■■■	hdiskpower4	■■■■/9A2	hdisk14 (alive)

IMPACT

Having fewer SAN I/O paths than required will result in reduced availability and performance. Moreover, in case there is only one SAN I/O path, a single point of failure exists between the server and the array which exposes this server to higher risk of downtime. In addition, applying load balancing algorithm such as round robin is not applicable when only path exists (performance can be significantly improved).

Business impact: HR

RESOLUTION

If the SLA deviations are by design, the resolution can be one of the following:

- Re-configure the SLA policy to be more accurate regarding the desired SLA
- Assign the SLA policy to sub-components under the host (e.g. VG, File system, Database service, etc.)
- Suppress the SLA deviations that are by design, using the Symptoms tab in the Tickets module

If the SLA deviations are not by design, re-configure the host or replicas so they will meet the SLA policy defined to them. Use the following section as a reference to fix the SLA issues in your environment:

The SLA policy "BC Tier 6 - Silver " is defined as followed:
Data Protection definition:
RPO:

- RPO 0 at any remote site (Enforced)

Retention:

- 2 remote physically protected copies with daily frequency (Enforced)

Availability definition:

- RAID level: RAID5
- Minimum number of SAN I/O paths: 2
- Minimum number of LUN maps: 2

The following table presents the SLA violation of AIX host ■■■■, grouped by the affected components:

Affected component	Affected business services	Violated SLA definition
VG ■■■■	* ■■■■ * ■■■■	minimum number of I/O paths (minimum 2 I/O paths)
VG ■■■■	* ■■■■ * ■■■■	minimum number of I/O paths (minimum 2 I/O paths)

TICKET 9

Gap ID	00305REPVIS	Name	Replication visibility
Severity	ERROR	Status	REOPEN
Categories	Extended Recovery Time		
Detected on	Mar 26, 2011 8:11:30 AM		
Verified on	Apr 13, 2011 12:56:55 PM		
Closed on	Nov 14, 2010 10:25:47 AM		

SUMMARY

Inconsistent access to replica Symmetrix devices of Veritas DG ■■■■ on 2 hosts (■■■■,■■■■)

DESCRIPTION

Veritas DG ■■■■ on 2 hosts (■■■■,■■■■) is stored on Symmetrix devices that have the following replica sets:

- Clone copy
- Clone-SRDF/S copy
- SRDF/S copy
- SRDF/S-BCV copy

One of these copies is visible by other hosts in an inconsistent way.

Typically, a target host does not necessarily need to access all existing replica sets of the Symmetrix device, but each replica set it does access - should be accessed completely (including all of its replica storage volumes).
In addition, all the replicas that it accesses should have the same layout path as each other.
Inconsistent access of the target hosts is usually a potential to extended recovery time or even inconsistent data on the target hosts.

The inconsistency was identified in the following replica set:
SRDF/S-BCV copy of Veritas DG ■■■■ on 2 hosts (■■■■,■■■■)

Physical volume	Source Symmetrix device	Replica Symmetrix device	Replica seen by
* emcpower138c on ■■■■ * emcpower119c on ■■■■	■■■■/68A	■■■■/F26	N/A
* emcpower21c on ■■■■ * emcpower4c on ■■■■	■■■■/4EA	■■■■/D82	Solaris host ■■■■
* emcpower34c on ■■■■ * emcpower23c on ■■■■	■■■■/5B2	■■■■/E56	Solaris host ■■■■
* emcpower41c on ■■■■ * emcpower26c on ■■■■	■■■■/4AE	■■■■/E9A	Solaris host ■■■■
* emcpower92c on ■■■■ * emcpower68c on ■■■■	■■■■/2FD	■■■■/1253	Solaris host ■■■■

Additional information to the table above:

- Replication layout path: SRDF/S-BCV

- Replica device are located at site:Vaguada DRTC
- Replica data age: Up to date
- Replica number of LUN maps: 2

IMPACT

Regarding inconsistent host access:

The storage volumes accessed by the target host are either:

- Incorrect (this may happen when it is not accessing the same number of replicas);
- An incomplete set of replicas;
- Are at a high risk of being inconsistent (meaning that it is accessing replicas with a different layout path).

This means that they may contain corrupt, outdated or incorrect data.

If these are the only remote replicas, then there is a high risk of data loss.

You should also check for the existence of gaps for the source Veritas DG.

Business impact: Billing

RESOLUTION

Make sure that target hosts are accessing the correct number of replicas, with the same layout path, and that they are of the same size and replication mode.

You may need to reconfigure both the source and the target.

If the target host does not need to access the replicas, dis-associate and unmap the devices from the host.

If the target host should access replicas of the source, then perform the following:

- First determine if replicas exist for each source storage volume. Allocate and configure missing replicas as necessary.
- Then configure access for the target host to the appropriate replica, verify that sufficient IO paths are allocated.

Use the table included in the ticket description section as a reference.

TOPOLOGY

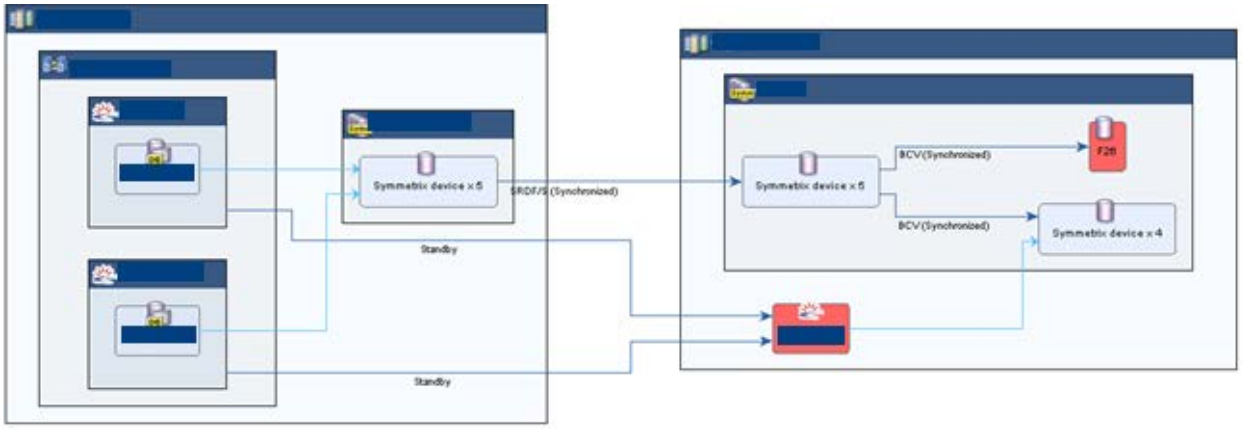
Symmetrix device x 5
SRDF/S (Synchronized)
Standby
Standby
Symmetrix device x 5
BCV (Synchronized)
BCV (Synchronized)
Symmetrix device x 4

TICKET 10

Gap ID	00300STRGSPLTGRP	Name	Storage group inconsistency
Severity	ERROR	Status	OPEN
Categories	Consistency		
Detected on	Dec 29, 2011 9:52:57 AM		
Verified on	Dec 31, 2011 12:58:07 PM		

SUMMARY

Device Group inconsistency was detected on replicas of the Symmetrix devices storing VG ■■■■ on AIX cluster node ■■■■ of cluster ■■■■ at site ■■■■

DESCRIPTION

Replica sets of VG ■■■■ on AIX cluster node ■■■■ of cluster ■■■■ at site ■■■■ are inconsistently connected to the Device Group.

This may adversely affect the hosts or VGs that use these replicas, since replica devices should typically be members of the same Device Group to insure data consistency.

The inconsistency was identified in 2 replica sets. The following tables provide more information on each set.

Replica set 1: SRDF/S copy of VG ■■■■ AIX cluster node ■■■■ of ■■■■

Physical volume	Source Symmetrix device	Replica Symmetrix device	Replica device groups
hdiskpower0	■■■■/60C	■■■■/60C	Device group ■■■■ on GNS server
hdiskpower122	■■■■/1C9	■■■■/1C9	N/A
hdiskpower1	■■■■/610	■■■■/610	Device group ■■■■ on GNS server

Additional information to the table above:

- Replication layout path: SRDF/S
- The replica devices are members of the RDF groups: RDF group 1 on Symmetrix ■■■■

Replica set 2: SRDF/S-BCV copy of VG ■■■■ on AIX cluster node ■■■■ of cluster ■■■■

Physical volume	Source Symmetrix device	Replica Symmetrix device	Replica device groups
hdiskpower1	■■■■/610	■■■■/D45	Device group ■■■■ on GNS server
hdiskpower122	■■■■/1C9	■■■■/15BA	N/A
hdiskpower0	■■■■/60C	■■■■/D41	Device group ■■■■ on GNS server

Additional information to the table above:

- Replication layout path: SRDF/S-BCV

Volume groups affected by this gap

- VG ■■■■ on AIX cluster node ■■■■ of cluster ■■■■

IMPACT

A Device Group inconsistency may indicate that replicas are split inconsistently, which might result in corruption, since the Symmetrix devices will fail to start or split simultaneously. This might result in data corruption or loss in the event of disaster.

Business impact: Document Management

RESOLUTION

Re-define the Device Groups so that they contain all source devices and all target devices of each affected replica.

TICKET 14

Gap ID	00304TRESTRCTMIS	**Name**	Partial replica sets
Severity	ERROR	**Status**	OPEN
Categories	Completeness		
Detected on	Dec 16, 2011 2:50:21 AM		
Verified on	Dec 17, 2011 3:14:54 AM		

SUMMARY

Control, temporary data, data, log and undo data files of 3 databases on 3 hosts (■■■■ ,■■■■ ,■■■■) are stored on Symmetrix devices that have partial replica sets

DESCRIPTION

Control, temporary data, data, log and undo data files of 3 databases on 3 hosts (■■■■ ,■■■■ ,■■■■) are stored on Symmetrix devices that have the following replica sets:

- 3 Clone copies
- SRDF/S copy
- SRDF/S-Clone copy

3 of these copies do not contain replicas of all source devices. This may represent data loss in the event of disaster (see impact).

The following table presents high level view of the replication structure

Physical volume	Source Symmetrix device	Replication layout (and # of paths)
*** emcpowerb on ■■■■** *** emcpowerb on ■■■■** *** emcpowerb on ■■■■**	**■■■■/A50**	**Clone (1)**
* emcpowero on ■■■■ * emcpowerp on ■■■■ * emcpowero on ■■■■	■■■■/AC0	* Clone (3) * SRDF/S (1) * SRDF/S-Clone (1)
* emcpowerp on ■■■■ * emcpowern on ■■■■ * emcpowerp on ■■■■	■■■■/AB0	* Clone (3) * SRDF/S (1) * SRDF/S-Clone (1)

Databases affected by this gap

- Control, data and log files of Oracle database ■■■■ on Linux host ■■■■
- Control, data and log files of Oracle database ■■■■ on Linux host ■■■■
- Control, data and log files of Oracle database ■■■■ on Linux host ■■■■

IMPACT

Incomplete copies of the database files exist, possibly resulting in an insufficient protection level and data loss upon disaster.

Alternatively, partial copies are sometimes just overly provisioned, or are leftovers from a data migration project, etc. In this case, the ticket represents a saving opportunity.

Business impact: Online Banking

RESOLUTION

Allocate and configure replicas for the devices specified to have none. Prefer configuring all the replicas (existing and new) on the same storage array, to ensure consistency of the replicas. If you are using storage consistency groups (such as CG, DG, RDF group, Consistent Group, Snapshot Session, etc.) to ensure data consistency, remember to include the new devices in the same storage consistency group, as required.

TOPOLOGY

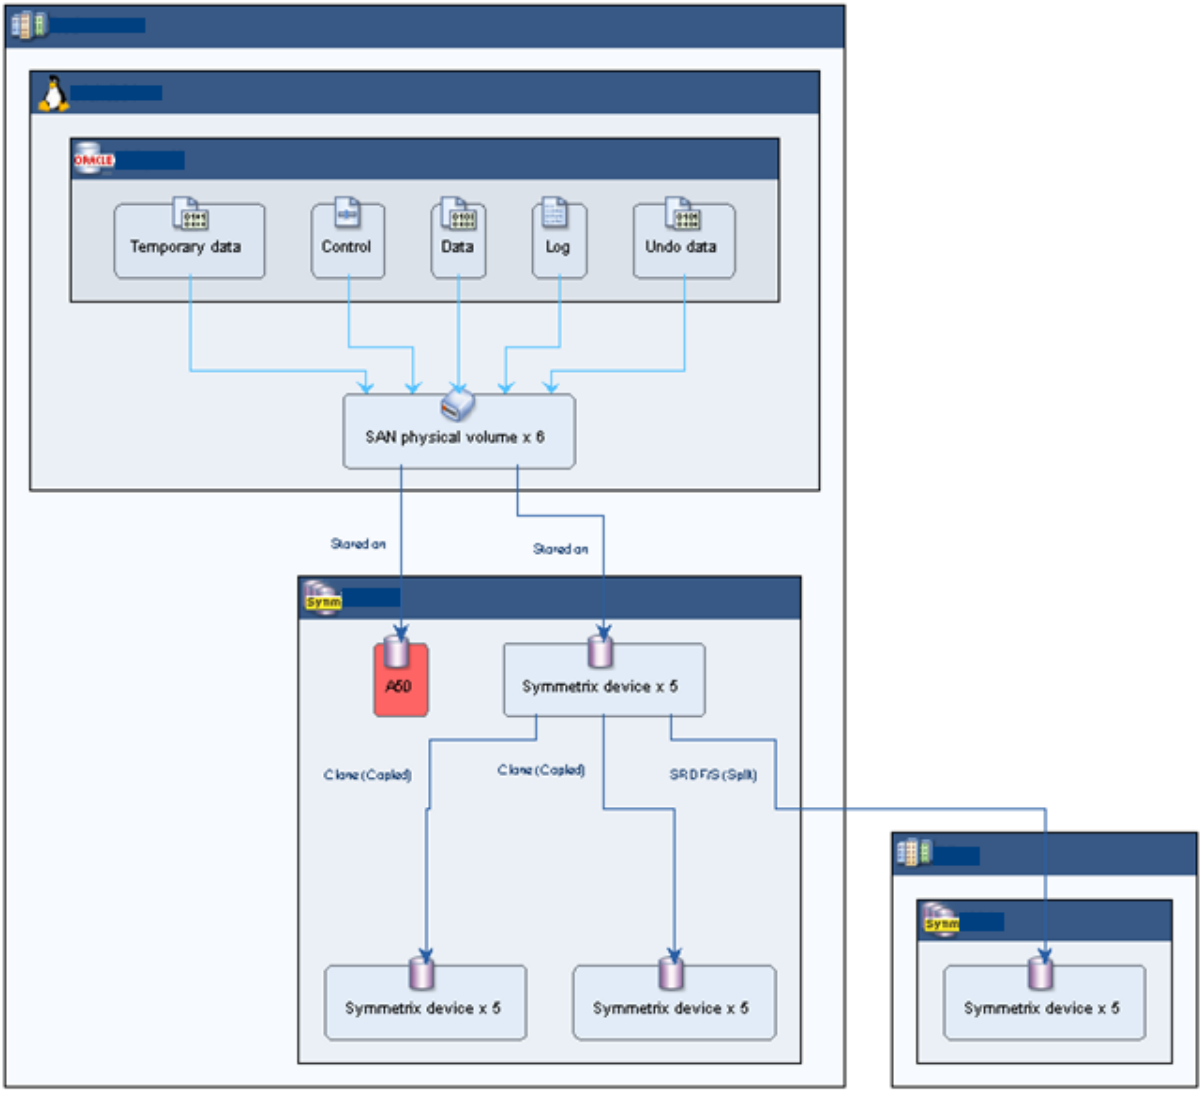

TICKET 17

Gap ID	00478VMDSUAC	**Name**	Unauthorized SAN access to VMware datastore
Severity	ERROR	**Status**	OPEN
Categories	Tampering		
Detected on	Dec 23, 2011 2:57:28 AM		
Verified on	Dec 24, 2011 2:36:06 AM		

SUMMARY

VMware datastore ■■■■ defined in datacenter ■■■■ is accessed through the SAN by unauthorized hosts

DESCRIPTION

VMware datastore ■■■■ defined in datacenter ■■■■ at site ■■■■ 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
■■■■/19FF	ESX host ■■■■	**EMC Fibre Channel Disk (naa.6000097000019260212553303139464 6)**
■■■■/19FF	ESX host ■■■■	**EMC Fibre Channel Disk (naa.6000097000019260212553303139464 6)**
■■■■/19FF	ESX host ■■■■	**EMC Fibre Channel Disk (naa.6000097000019260212553303139464 6)**

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.

Business impact: CRM

RESOLUTION

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

TOPOLOGY

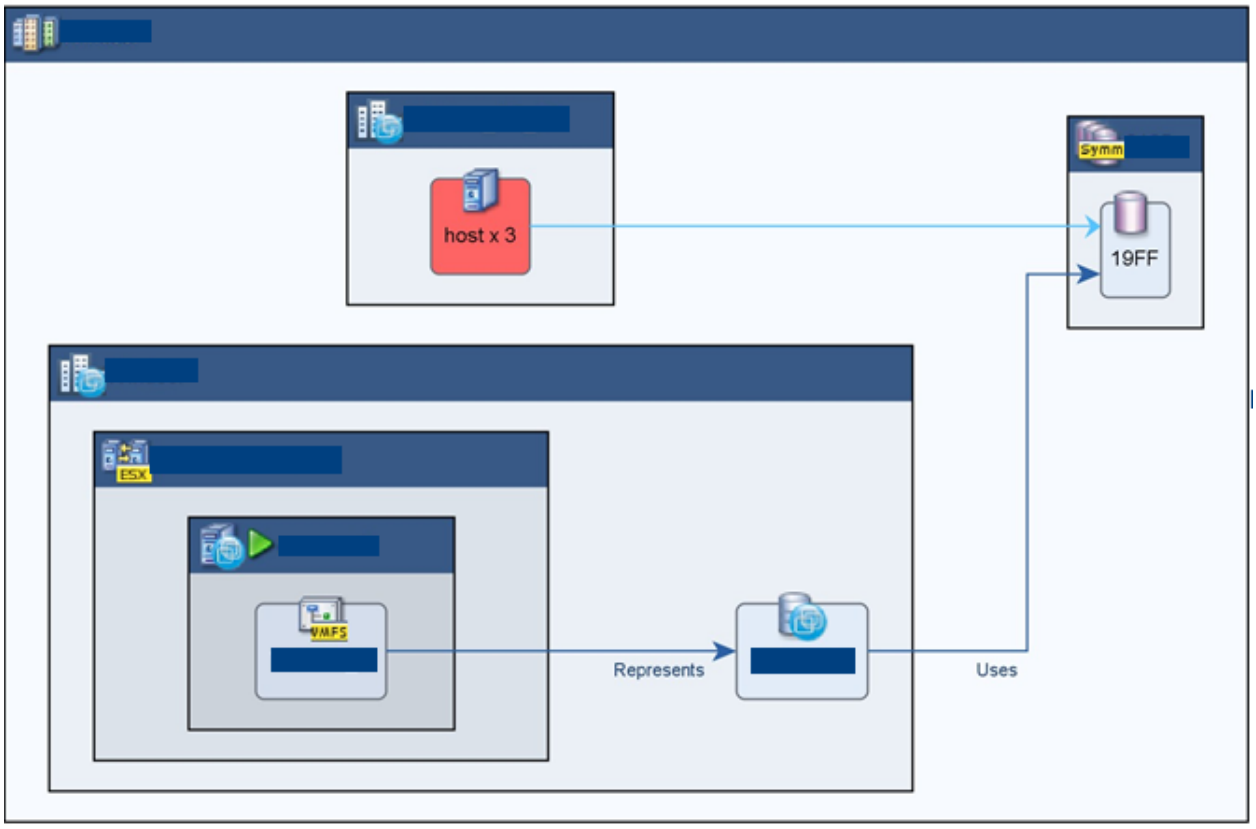

TICKET 18

Gap ID	00318SANIOSHSTG	**Name**	Hosts accessing shared storage with different number of SAN I/O paths
Severity	ERROR	**Status**	OPEN
Categories	Performance, Availability		
Detected on	Dec 03, 2011 10:06:22 PM		
Verified on	Dec 14, 2011 9:47:58 PM		

SUMMARY

The shared HDS logical units storing datastore ■■■■ of datacenter ■■■■ at site **DR** are accessed with different number of SAN I/O paths

DESCRIPTION

The shared HDS logical units storing datastore ■■■■ of datacenter ■■■■ are accessed by 11 nodes of the ESX cluster ■■■■. These cluster nodes are accessing these HDS logical units with different number of SAN I/O paths. This might result in performance degradation and increased risk of unexpected downtime (see impact).

Datastore ■■■■ of datacenter ■■■■

Physical volume	Source HDS logical unit	Number of paths from ■■■■	Number of paths from ■■■■	Number of paths from ■■■■	Number of paths from ■■■■
HITACHI Fibre Channel Disk (naa.60058a90) on 11 hosts	■■■■/0063	2	1	2	1

Datastores affected by this gap

- Datastore ■■■■ of datacenter ■■■■

File systems affected by this gap

- Local file system **C:** on Windows VM ■■■■

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.

Business impact: Trading

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 20

Gap ID	00397HBASO	**Name**	Suboptimal HBA configuration
Severity	WARNING	**Status**	OPEN
Categories	Performance, Availability		
Detected on	Dec 07, 2011 9:46:33 PM		
Verified on	Dec 14, 2011 9:49:27 PM		

SUMMARY

I/O redundancy has a single-point-of-failure in the HBA level on Windows host ■■■■ at site ■■■■

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 Windows host ■■■■ at site ■■■■ are configured with multipathing, 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 local file system C:\LUN0\ on Windows host ■■■■ at site ■■■■:

Host	Pseudo PV	HDS logical unit	HBA port	HBA serial
Windows host ■■■■	**Disk6**	**■■■■/1291**	**\\.\Scsi2:**	0000c92e5426

The following list displays the available HBAs on ■■■■:

- 0000c92e40b4/**\\.\Scsi3:**
- 0000c92e5426/**\\.\Scsi2:**

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.

Business impact: Email

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 23

Gap ID	00255CLHSBP	**Name**	CLARiiON Hot Spare best practice violation
Severity	WARNING	**Status**	OPEN
Categories	Redundancy, Best Practice		
Detected on	Dec 27, 2011 6:29:05 AM		
Verified on	Dec 01, 2011 4:22:13 AM		

SUMMARY

CLARiiON ■■■■ at ■■■■ is violating hot spare best practice

DESCRIPTION

Hot spares are disks that are pre-allocated at the time of configuration. When a disk fails, the hot spare is switched into operation and replaces it. According to EMC best practice it is recommended have a minimum of 1 spare drive for up to 30 drives (3.33%) (per drive type). In addition, the hot spare must be at least the same size as the largest size disk (of the same device type) in the array, and it is highly recommended that the hot spare speed will be at least the same speed as the fastest disk (of the same device type) in the array.

A gap has been detected where CLARiiON ■■■■ at **unknown site** is violating some of the best practices mentioned above.

The following table presents the CLARiiON's hot spare violations:

Drive type	Violated Best Practices	# of allocated physical disks	Actual/Recommended # of hot spares	Fastest disk	Largest disk	Hot Spares
Fibre Channel	Insufficient number of hot spares	250	5/9	**4Gbps**	**402.61GB**	* **Bus 0 Enclosure 0 Disk 14** (402.61GB, 4Gbps) * **Bus 1 Enclosure 0 Disk 14** (268.4GB, 4Gbps) * **Bus 7 Enclosure 1 Disk 14** (268.4GB, 4Gbps) * **Bus 5 Enclosure 0 Disk 14** (268.4GB, 2Gbps) * **Bus 5 Enclosure 1 Disk 14** (268.4GB, 2Gbps)

IMPACT

Upon disk error, automatic correction actions taken by the array may fail or work sub-optimally due to insufficient number of spare disks, or insufficient disk space.

Business impact: Email

RESOLUTION

There should be at least one hot spare disk for every 30 drives on the CLARiiON storage system. It is up to the customer to configure more hot spares. For ease of management, it is also recommended that the hot spare be configured on the last drive slot on a disk-array enclosure/shelf. However, the hot spare may be configured anywhere in the system with exception of the vault drives that are used for cache vault and certain other internal purposes. The vault drives are the first five drives on the CX series. Their location varies for prior generation of products.

In addition to the number of hot spares allocated, make sure that for each device type, the hot spare's size is equal or larger than the largest disk on the array. Also make sure that the hot spare speed is equal or faster than the fastest disk on the array.

TICKET 27

Gap ID	00832DGHBF	**Name**	DataGuard heartbeat failure
Severity	ERROR	**Status**	REOPEN
Categories	Replication SLA, Completeness		
Detected on	Oct 27, 2011 6:35:36 AM		
Verified on	Dec 01, 2011 4:28:14 AM		
Suppressed on	Dec 03, 2011 4:33:49 AM		

SUMMARY

DataGuard heartbeat failures were detected on primary Oracle database ■■■■ on Linux host ■■■■ at site ■■■■

DESCRIPTION

DataGuard heartbeat samples the communication between the primary and the standby databases. Heartbeat failure usually indicates communication problems between the databases. When communication problems exist, the log transfer between databases is halted, resulting in a period of time in which the standby database is not synchronized with the primary database.

The following table presents the heartbeat failures detected on primary Oracle database ■■■■ on Linux host ■■■■:

Standby database	Date	Error Code	Heartbeat failure start time	Heartbeat failure end time	Duration
HOST=■■■■; PORT=1522;SERVICE_NAME=■■■■;INSTANCE_NAME=■■■■	Oct 30, 2011	12541	3:48:02 PM	4:48:02 PM	60:00 minutes
HOST=■■■■; PORT=1522;SERVICE_NAME=■■■■;INSTANCE_NAME=■■■■	Oct 08, 2011	12560	11:34:18 PM	12:04:37 AM	30 minutes

Error code description:

- ORA-12541: TNS:no listener
- ORA-12560: TNS:protocol adapter error

IMPACT

Since the log transferring between the databases is halted and the standby database is not synchronized with the primary database - failover during heartbeat communication problem may lead to data loss and RPO SLA breach.

Business impact: Data Warehouse

RESOLUTION

Check any communication issues between the primary and the standby databases.