Business continuity with IBM HyperSwap, Oracle Real Application Cluster (RAC), and VMware Site Recovery Manager

A solution overview

Overview

Challenge

Do you have a business continuity plan to mitigate threats and risks for your business in the event of potential disaster? Do you have an orchestrated way to quickly get your business running post disaster?

Solution

Consider a multifold solution that brings together the best of the industry technologies together. IBM HyperSwap provides an active/active replication of volumes across metro distances, Oracle Real Application Cluster provides good scalable application environment, and VMware Site Recovery Manager provides an environment complete with automated orchestration and disaster recovery capabilities.

In today's world, coming up with a business continuity plan (BCP) is imperative to the sustainability of your business. Without a well-thought-out plan in place, it is highly unlikely that your company will be able to survive and recover from disasters.

The BCP involves creation of a strategy through the recognition of threats and risks facing a company, with an eye to ensure that personnel and assets are protected and are able to function in the event of a disaster.

Business continuity is a proactive plan to avoid and mitigate risks associated with a disruption of operations. It details the steps to be taken before, during, and after an event to maintain the financial viability of an organization. People often mistake disaster recovery as business continuity plan. It is in fact a reactive plan for responding after an event.

However important the BCP is, there are several major roadblocks (such as business priority, prohibitive costs, high complexity, and willingness of the people involved) for successful implementation.

Having a BCP enhances an organization's image with employees, shareholders, and customers by demonstrating a proactive attitude. Additional benefits include improvement in overall organizational efficiency and identifying the relationship of assets and human and financial resources to critical services and deliverables.

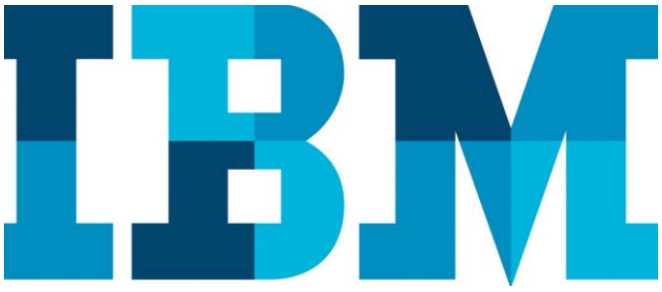

The BCP help mitigate risks from potential disasters that include:

- Natural disasters
- Accidents / Sabotage / Cyber attacks
- Power and energy disruptions
- Communications, transportation, safety, and service sector failure

Although there are several solutions to build a business continuity plan, this paper restricts discussion to the following three technologies:

- IBM® HyperSwap®
- Oracle Real Application Cluster (Oracle RAC)
- VMware Site Recovery Manager

These technologies play a vital role within the technology stack such as storage, application, and infrastructure.

Assumptions

The proposed solution includes the following assumptions:

- Good understanding about setup and working knowledge of Oracle RAC. A functional two node Oracle RAC is available at any one of the sites. Refer to Oracle RAC installation guide at: https://docs.oracle.com/database/122/RILIN/title.htm
- Good understanding about setup and working knowledge of VMware Site Recovery Manager. Refer to VMware Site Recovery Manager installation guide at: https://docs.vmware.com/en/Site-Recovery-Manager/6.5/com.vmware.srm.install_config.doc/GUID-723EAC1B-AC21-4CAA-9867-627CA8CB680A.html
- Availability of two IBM FlashSystem® A9000 or A9000R systems connected using Fiber Channel (FC) connectivity with software version 12.1. For more details, refer to the documentation at: https://www.ibm.com/support/knowledgecenter/en/STKMQV_7.8.0/com.ibm.storage.vflashsystem9000.7.8.0.doc/svc_hyperswap_configuration.html

About IBM HyperSwap

IBM storage solutions provide the speed and performance required for structured or unstructured data. IBM FlashSystem A9000 and A9000R storage systems are engineered to meet modern high-performance storage requirements with ultra-low latency, cost-effectiveness, operational efficiency, and mission-critical reliability.

IBM FlashSystem A9000 and A9000R storage systems, both existing and new, can be clustered to deliver active-active all-flash storage access, using the new HyperSwap capability.

HyperSwap is made of 100% built-in system software, and requires no extra license or hardware. After the existing FlashSystem A9000 and A9000R systems are live upgraded to version 12.1, or when the new systems are deployed, they are HyperSwap ready.

HyperSwap enables enterprise IT systems to autonomously fail over within seconds, without human intervention and without mind-boggling automation masterpieces. HyperSwap uses a metro-distance FlashSystem A9000/A9000R stretch-cluster to protect against loss of access to mission-critical data. HyperSwap distances are limited by synchronous mirroring distances (typically up to 100 km apart) depending on the available bandwidth and impact of latency on the host. If the application is latency sensitive, you might want to further limit the distance to a maximum of 75 km.

For more details on storage systems connectivity refer to section 2.3.1 from IBM Redpaper™: redbooks.ibm.com/redpapers/pdfs/redp5434.pdf

The HyperSwap function is available on IBM storage technologies that includes IBM System Storage® DS8880F storage family, IBM FlashSystem A9000/A9000R storage family, and IBM Storwize® family that includes IBM Spectrum Virtualize™ software.

Figure 1 shows an overview of HyperSwap overview within an application ecosystem.

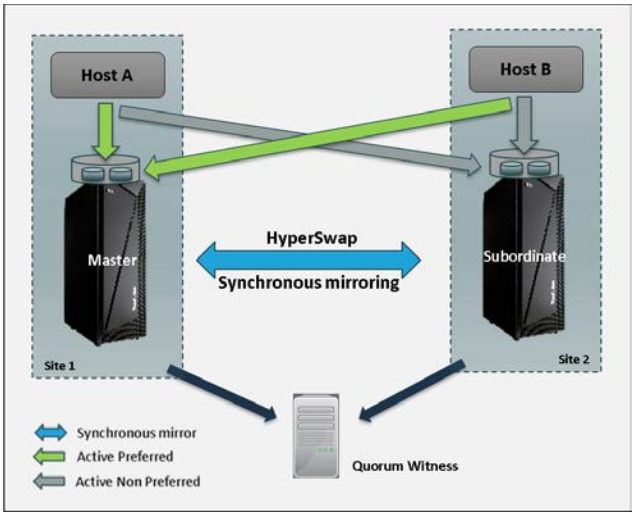

Figure 1: HyperSwap overview

The paper focuses on IBM FlashSystem A9000 and A9000R with IBM HyperSwap feature for the proposed solution.

About Oracle Real Application Cluster

Oracle RAC feature gained a lot of popularity with Oracle release 10*g*. Oracle RAC helps eliminate database server as single point of failure (SPOF) within an application ecosystem.

It focuses on recoverability as any other surviving node in the cluster can recover the unfinished transactions in an event of node failure.

The cluster software error detection mechanism detects the failure of a resource within the cluster and attempts to restart or relocate the failed resource to other nodes and lastly, scalability is the most useful feature in terms of dynamic expansion or reduction of the cluster.

The combination of Oracle RAC and Oracle Data Guard has proven itself as one of the most compelling solutions for scalability, high availability, and disaster recovery.

For the proposed solution, the test team refers to the Oracle RAC as application layer.

About VMware Site Recovery Manager

VMware vCenter Site Recovery Manager is a disaster recovery management solution. Site Recovery Manager helps in business continuity and disaster recovery solution helps you to plan, test, and run the recovery of virtual machines between a protected VMware vCenter Server site and a recovery vCenter Server site.

You can implement Site Recovery Manager for different types of recoveries from the protected site to the recovery site.

Planned migration is the orderly evacuation of virtual machines from the protected site to the recovery site. Planned migration prevents data loss when migrating workloads in an orderly fashion. For planned migration to succeed, both sites must be running and fully functioning.

Disaster recovery does not require both protected and recovery sites be up and running even if the protected site goes offline unexpectedly. During a disaster recovery operation, failure of operations on the protected site is reported but is otherwise ignored.

Site Recovery Manager orchestrates the recovery process with the replication mechanisms to minimize data loss and system down time.

IBM FlashSystem A9000 and A9000R HyperSwap feature is certified with VMware Site Recovery Manager as well as VMware stretch clustering

without Site Recovery Manager. Refer to the "Get more information" section for more information on this.

For the proposed solution, the test team refers to VMware Site Recovery Manager as infrastructure layer.

Solution overview

This section provides an overview of the proposed solution bringing together all the various infrastructure layers mentioned so far.

Architecture

Software

- CentOS Linux 7.2
- Oracle 12c Release 2 database and grid infrastructure software
- VMware Site Recovery Manager 6.5
- Swingbench load generator

Hardware

- Intel x86 processor-based systems
- Brocade SAN Fabric
- VMware suite of products
- IBM FlashSystem® A9000 and A9000R

Network/SAN

- 16Gb FC cards and switches
- 10Gb Ethernet cards and switches

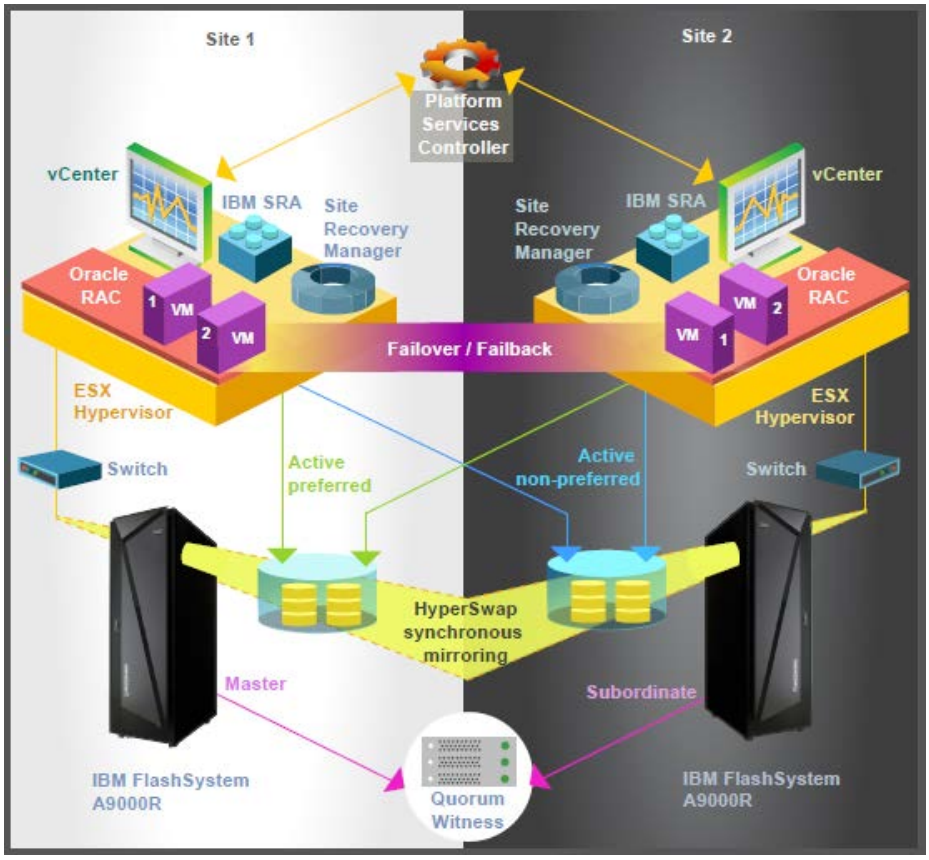

Figure 2: Integrating IBM HyperSwap, Oracle RAC, and VMware Site Recovery Manager

When building this solution, start with the foundation layer of storage. In this example, there are two sites at metro distance connected using Fibre Channel (FC) connectivity for HyperSwap configuration. On both sites, IBM FlashSystem A9000 or A9000R storage systems are deployed.

First, consistency groups are created on IBM FlashSystem A9000 systems and they are configured in a HyperSwap relationship. The required volumes are created on the IBM FlashSystem A9000 systems and HyperSwap relationship is assigned to them. The volumes are then moved to the consistency group and are

mapped to the ESXi servers on both sites, which are used by ESXi servers in active/preferred and active/non-preferred configurations.

Note that the HyperSwap relationship is automatically activated as soon as it is created.

For more details on HyperSwap, refer to IBM HyperSwap for IBM FlashSystem A9000 and A9000R in IBM Redbooks®.

After the storage configuration is in place, you need to create the VMware infrastructure.

The team used two VMware vCenter servers deployed across two sites but pointing to the common VMware Platform Services Controller (PSC). Storage policies were created on both VMware vCenter servers using custom-built tags to associate datastores created on IBM FlashSystem A9000 HyperSwap volumes. The datastores created are further used to store virtual machines.

After installing VMware Site Recovery Manager instances on both the sites, IBM Spectrum Accelerate™ Family Storage Replication Adapter (SRA) is installed on the Site Recovery Manager server instances on both the sites. IBM Spectrum Accelerate Family Storage Replication Adapter can be downloaded from the VMware Site Recovery Manager website.

IBM FlashSystem A9000 systems are then configured to VMware Site Recovery Manager across both the sites. This solution uses storage policy based protection groups which are tagged to datastores created on IBM FlashSystem A9000 storage HyperSwap volumes in Site Recovery Manager. A recovery plan is created using the storage policy-based protection group. After the recovery plan is created, it is a good practice to test the recovery plan by performing a failover or failback of the virtual machines from both sites to make sure that Site Recovery Manager is correctly configured and is working fine.

After the VMware infrastructure was put in place, the team converted the virtual machines on respective sites as Oracle RAC nodes. In this solution, the team created two Oracle RACs on each site consisting of virtual machines available locally. The VMs on the respective sites share the data.

The benefit of this configuration seen in the event of failover scenarios is described in Table 1.

No	Scenario	Result
1	System A full failure	• System B is the master
2	System A partial failure	• System B is the master • System A local volume is unavailable for I/O
3	A<->B connectivity failure	• System A is the master • System B local volume is unavailable for I/O
4	System B failure	• System A remains as the master
5	Quorum witness failure	• High availability is down

Table 1: Failover scenarios and expected results

In each case, based on the severity of the situation, administrators can decide to perform a planned or unplanned migration of the VMs from one site to another.

It does not have to be a disaster all the time in order to test the setup. So, during planned failovers, the solution will work just as expected.

Summary

The main focus of this paper is on BCP in the event of a disaster. This paper showcases the solution by combining some of the best features from industry-recognized applications and solutions. This solution paper summarizes IBM FlashSystem A9000 and A9000R HyperSwap configuration to achieve an active/active storage configuration. On top of this, VMware Site Recovery Manager provides orchestrated planned or unplanned migration of virtual machines to provide transparent failover of workload such as Oracle RAC across the sites.

Acknowledgement

Special thanks to the team members who were directly or indirectly involved in this project. The author wishes to thank the IBM FlashSystems A9000 development team and Ran Harel for their valuable input during HyperSwap configuration.

Get more information

The following websites provide useful references to supplement the information contained in this paper:

- IBM FlashSystem A9000 and A9000R HyperSwap Resources
  https://developer.ibm.com/storage/2017/06/09/flashsystem-a9000-a9000r-hyperswap-available/

- IBM HyperSwap for IBM FlashSystem A9000 and A9000R Redpaper
  https://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/redp5434.html?Open

- IBM Spectrum Accelerate Family Storage Replication Adapter
  https://marketplace.vmware.com/vsx/solutions/ibm-xiv-storage-replication-adapter

- VMware knowledge base article on vSphere Stretch Cluster Solutions with IBM FlashSystem A9000 and A9000R HyperSwap
  https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2151070

About the authors

Shashank Shingornikar is an ISV solutions engineer in IBM Systems. Shashank has experience in Oracle technology products, Oracle data migration, and disaster recovery solutions. Currently, he is working on enabling Oracle solutions for IBM storage systems.

You can reach Shashank Shingornikar at shashank.shingornikar@in.ibm.com

Mandar Vaidya is an ISV solutions engineer in IBM Systems. Mandar has been working to enable VMware solutions for IBM storage systems.

You can reach Mandar Vaidya at mandar.vaidya@in.ibm.com

Yossi Siles, a senior offering manager, joined IBM in 2010. Yossi drives FlashSystem A9000, A9000R and IBM XIV® Storage System software from roadmap inception through definition, to delivery. Before joining IBM, Yossi was a senior program manager at Microsoft, focusing on network-perimeter security.

You can reach **Yossi Siles** at SYOSSI@il.ibm.com

IBM Systems
3039 Cornwallis Road
RTP, NC 27709

Produced in the United States of America