IT Disaster Recovery and Business Continuity Based On VMware SRM Solution for Kuwait Oil Company (KOC)

Mohammad Matar Alshammari
Department of Computer Science, Kulliyyah of Information and Communication Technology, International Islamic University Malaysia, Kuala Lumpur, Malaysia
mdshammari@gmail.com

Ali Amer Alwan
Department of Computer Science, Kulliyyah of Information and Communication Technology, International Islamic University Malaysia, Kuala Lumpur, Malaysia
aliamer@iium.edu.my

Abstract—The **IT center of Kuwait Oil Company (KOC) is one of the largest information centers in the State of Kuwait. It provides many services to KOC departments and other Kuwait oil sector companies (K-Companies). These services need to ensure High Availability (HA) and flexibility in their operations. The aim of this paper is to investigate the impact of the damage that affects the Business Continuity (BC) and data availability during and after disasters. The paper presents a VMware Site Recovery Manager (SRM) as a Disaster Recovery (DR) and BC solution with open-source to insure HA of IT services at KOC. Moreover, the propose solution introduces a better understanding of Disaster Recovery Plan (DRP) and Business Continuity Plan (BCP) methodology to enhance the KOC IT center by creating and validating a plan for maintaining continuous business operations before, during, and after disasters and disruptive events.**

Keywords- Kuwait Oil Company; Business Continuity Plan; Disaster Recovery Plan; High Availability; VMware; Site Recovery Manager.

I. Introduction

In today's business environment IT resources involving data are among an organization's most important assets. Natural disasters such as earthquakes, cyclones, hurricanes, and floods can destroy these assets. They can also be intentionally destroyed by computer viruses, hackers, vandalism, and terrorist attacks.

In view of this, organizations need to be prepared and equipped against any disaster in order to maintain their survival and reputation by quickly recovering the data and continuing their operations. Therefore, effective and efficient disaster recovery and business continuity plans are not optional, but critical for an organization's success.

The Disaster Recovery (DR) and Business Continuity Plan (BCP) literature offers prescriptive procedures for planning and implementing DR/BCP. Documents and papers are presented along with case studies that raise a specific awareness. Although, many works have been published in recent years, still it is challenging to acquire relevant information concerning regional events such as Oman's Cyclone Guno [1], due to the lack of an appropriate design to suit the available facilities in their environments to organizations in different areas. In [2], differences between disaster and crisis are introduced with extensive definitions, models, characterizations, criteria and types of disasters and crises.

The requirements for an inclusive, cohesive, integrated, and multi-dimensional disaster policy are presented in [3]. The paper reports both durability and impairment of alternative perspectives on the disasters. It also suggests that the broad concept of vulnerability is best suited for assimilating academic work and facilitating policy guidance for professionals in the business. Moreover, the paper recommends several issues for determining responsibilities and abilities, reducing risks, raising resistibility and flexibility. In [4], business continuity is performed through the duplication of hardware infrastructures, data replicating, and immediate applications availability. Realizing the problems caused by in consistencies, and quality control measures for emergency plans are proposed in [5] by stating 18 aspects that can be standardized across emergency plans for BC and DR. The emergency plan concentrates on local bodies and provides guidelines for testing, use, and revision.

Disaster preparedness plans and BCP are integrated into one model in [6]. Disaster preparedness plans involve preparedness, response and recovery, while business continuity plans consists of response, stabilization, assessment and business continuity. In [7], the outcomes from a survey of 274 executives in India are assessed by measuring the impact of computer disasters on the management of information. The results show that the information in all relevant companies is negatively influenced due to the occurrence of such disasters. In addition, hardware failures and software viruses are disasters that take place in most organizations.

The influence of the Internet and communication during disasters are discussed in [8]. The paper mentions the value of exploiting effective internet-enabled communication such as text messages over Personal Digital Assistants (PDAs) or cellular phones, instant messaging, web pages, and emails

978-1-5090-4580-8/16/$31.00 ©2016 IEEE

during a disaster. In [9], the standing of Critical Success Factors (CSFs) to implement a BC/DR program that have been replaced in earlier research is highlighted, especially after 9/11. The literature states several CSFs that are not indicated in previous studies. Whereas in [10], establishing a connection between emergency response plans, DR and BCP, was suggested as an approach to effective management of crisis.

The role of the Indian insurance industry in managing disasters is introduced in [11]. It states that in several developing countries, most of the losses suffered in natural disasters are not insured because of the lack of a proper plan to handle the disaster. This situation arises due to a lack of appropriate policies, limited purchasing power, and scant interest in insurance. The paper offers an integrative approach for managing disasters.

The DR levels are classified through those two important measures: Recovery Time Objective (RTO) and Recovery Point Objective (RPO), which are the primary goals that should be achieved anytime, assessing a better choice in a given operating and capital costs [12] [13]. RPO defines the data that we might afford to lose, the lower the RPO the higher the total expense of preserving the actual infrastructure environment for recovery [14]. This is further investigated by Lenk and Tai [15] who stated that, RPO indicates the highest possible appropriate period in between a couple of backups while RTO describes the most acceptable time period a business process might be disrupted as indicated in Figure 1.

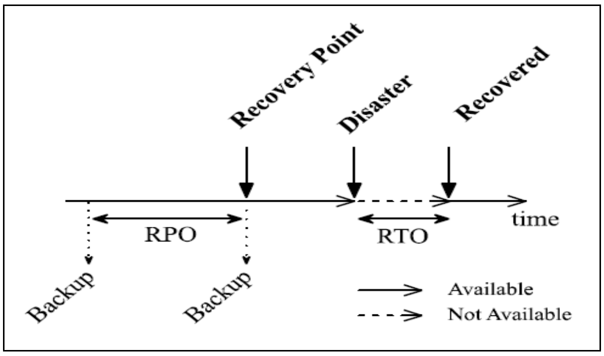

Figure 1. Recovery Point Objective (RPO) and Recovery Time Objective (RTO)

II. Problem Description

The IT center of KOC is one of the largest information centers in the State of Kuwait. The center provides many services to KOC departments and other Kuwait Oil Sector Companies (K-Companies). The growing number of services gives rise to the following drawbacks:

- The DRP is not prepared correctly.
- Some systems such as the Gas Management Information System (GMIS) and the Hospital Information System (HIS) are not in duplicated at the recovery site.
- Databases are not synchronized between the production and recovery sites.
- The recovery site is activated manually, which is time consuming.
- BCP and Business Impact Analysis (BIA) do not exist.

This paper attempts to remedy the aforementioned drawbacks by providing a disaster recovery solution that will assist IT services at KOC. It also introduces a better understanding of Disaster Recovery Plan (DRP) and Business Continuity Plan (BCP) methodologies in order to enhance KOC by creating and validating a plan for continuously maintaining the business operations before, during, and after natural disasters and intentionally disruptive events.

III. Business continuity and Disaster recovery

A. *Business Continuity Plan*

A Business Continuity Plan (BCP) is a procedure used to establish and validate a plan for continuously maintaining the operations of a business before, during, and after disasters and disruptive events. It considers the management of the operational components that allow a business to function normally in order to create revenue. BCP guarantees the shift of critical systems to another environment while the originals are being repaired. Also, it sets the right people in the right places and performs business in different modes by involving the shareholders through different stations until everything returns to normal.

In addition, BCP enhances organizations to be able to continue their operations regardless of the nature of a potential disruption by providing guidance to IT staff to follow an emergency plan in order to recover and resume IT services when operations are unexpectedly disrupted [16]. Specifically, preplanned procedures allow an organization to successfully do the following:

- Producing an immediate and valid response to emergency events.
- Ensuring safety and protecting lives.
- Reducing business impact.
- Resuming the services of critical business.
- Working with vendors during the recovery time.
- Reducing the confusion during a disaster.
- Ensuring continuity of business functions.
- Obtain “up and running” quickly after a disaster.

These procedures must be preplanned in such a way that ensuring timely and orderly resumption of an organization’s business cycle, at the same time can be executed without interruption or minimal to time-sensitive IT service operations.

B. *Disaster Recovery Plan(DRP)*

Disaster is defined as a destructive or debilitating occurrence that compromises the operational availability of the system for an unacceptable period of time. Disasters are different from general failure in both severity and degree of impact. System failures don’t necessary impair system capability. Disasters that cannot be ameliorated by existing failure prevention systems are generally the result of catastrophic events including, but not limited to, human intervention, severe weather, floods or fire. Since a disaster destroys and interrupts the continuity of business operations within an IT center, the response requires usage of additional infrastructures.

DRP is a part of BCP and deal with the immediate impact of the event. Recovering from a server outage, security violation, or hurricane, all fall into this category. Usually, DRP consists of several forethought steps which are prepared to be applied in planning stages. The implementation of these steps is swift when a disaster occurs even though the situation during the disaster is almost never exactly as planned. In this direction, the resources can be controlled in accordance with the prepared steps. Presently, DRP quickly indicates the influence of a disaster and addresses the immediate results.

IV. Establishing High Availability

A. High Availability

High Availability (HA) is the first crucial step in maintaining business continuity of IT services in case of failures or problems that are not caused by major disasters and can be managed locally in accordance to ordinary operational procedures without resorting to DRP. One of the earliest promises of virtualization is the ability to preserve the virtualized systems online and operate in spite of existing problems underlying hardware by permitting a Virtual Machine (VM) to be run in any host virtual environment. HA is a design methodology exploited to guarantee the availability and uptime of VMs.

In general, there are two types of downtimes mitigated via HA that provides using virtualization technologies:

- Planned downtime: It is a time for scheduling the maintenance and upgrading during in which a system cannot be used for normal productive operations.
- Unplanned downtime: It is a time in which a system cannot be used for normal productive operations due to unforeseen failure in hardware/ software components or operator mistakes.

B. Virtualization

Virtualization is a technique for simultaneously running multiple operating systems on a single computer and making one computer operate as multiple computers as shown in Figure 2 It is considered a framework/methodology to divide computer hardware resources into multi-execution environments via implementing one or more concepts/technologies, such as time-sharing, hardware and software partitioning, partial/complete machine simulation,, quality of service, emulation, and so on.

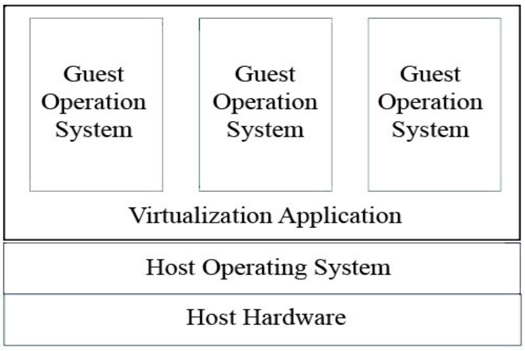

Figure 2. Virtualization environment.

C. VMware Site Recovery Manager

VMware Site Recovery Manager (SRM) is a disaster recovery and business continuity open-source solution that assists IT organizations in planning, testing, and executing an emergency failover or scheduled migration of datacenter services from one site to another. SRM is supported by VMware vCenter to provide integration with array based replication, discovery and management of replicated data stores, and automated migration of inventory from one vCenter to another.

SRM server is used to coordinate the operations of replicated storage arrays and vCenter servers at two sites. This implies that as VMs at the production site are shut down, the VMs at the recovery site start up and use the data replicated from the production site to assume responsibility for providing the same services. Transfer of services from one site to the other is controlled by a recovery plan that specifies the order in which VMs are shut down and started up, the computer resources they are allocated, and the networks they can access. In Figure 3, SRM allows testing of a recovery plan by using a temporary copy of the replicated data. The process does not disrupt ongoing operations at either site.

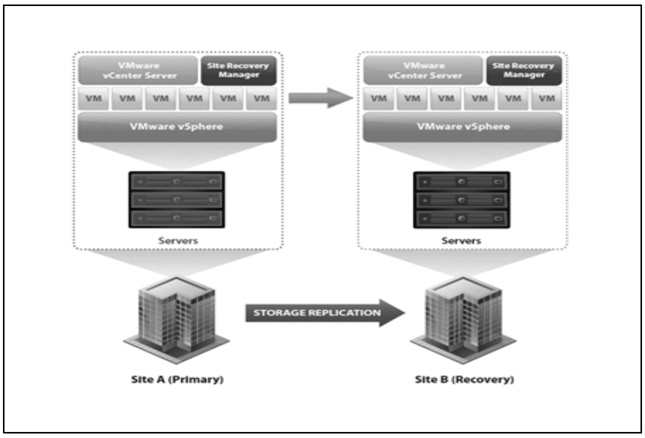

Figure 3. SRM production and recovery sites.

V. Proposed Solution for IT Disaster Recovery in KOC

Site Recovery Manager (SRM) is an end-to-end DR automation open-source solution. SRM automatically customizes VMs in a way that they can run at a recovery site. It is designed to protect VMs residing in a datastore on replication storage. In the event of a storage array failure or complete site failure VMs can be transferred over to a remote datacenter.

The proposed solution involves the implementation of SRM with a VM environment for IT services at KOC by providing references to install and configure SRM within a test environment at KOC. It also introduces a framework to implement a DRP at the recovery site. Setup, testing, evaluation and failover are taken into account as well.

The solution assumes two sites: production and recovery. In this particular assumption, the production site provides business-critical data store services whereas the recovery site acts as an alternative facility to which these services can be migrated. The production site is located in the main computer center of KOC in which a virtual infrastructure supports a critical business need. The recovery site is in the main KOC office, which is 1.5 miles from production site. The logical architecture of SRM in KOC is shown in Figure 4. Obviously, SRM needs several requirements for vCenter configurations at each site as follows:

- Each site must include at least one vCenter datacenter.
- The recovery site supports array-based replication with the production site, and has hardware and network resources to support the same VMs and workloads as the production site.
- One VM must be located on a replicated datastore at the production site. The datastore must be supported by a storage array that is compatible with SRM.
- The production and recovery sites connect via fiber channel network.
- The recovery site has the ability to access the same public and private networks as the production site, and is not necessarily in the same range of network addresses.

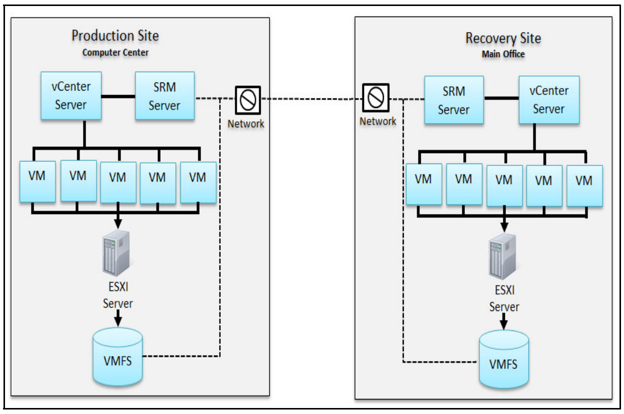

Figure 4. SRM logical architecture in KOC.

Therefore, the objectives of the proposed solution are to provide guidance to the IT-services group at KOC in order to:

- Describe disaster scenarios that may affect IT services at KOC and high availably that can be replicated to remote recovery sites for evaluation.
- Provide suitable integration of VMware SRM in the IT department at KOC to react to disasters and manage crisis situations.
- Identify the ability to simulate and test failover processes without impacting existing operations to insure the processes will be performed correctly during a real disaster.
- Demonstrate the process of restoring normal operation of the original production site after a failover.

A. *Setup of Proposed Solution*

Figure 5 highlights the configuration of an environmental test of the proposed solution in KOC. It requires two servers: production and recovery. Both of them include six VMs in one ESXi host.

The implementation of an environmental test uses SRM prerequisite to perform the following tasks:

- Install a VMware ESXi host at both sites to run multiple VMs.
- Create six VMs at each site.
- At both sites, set up and configure the applicable networking.
- Ensure each site consists of a VMware vCenter Server and a SRM plug-in in two VMs.
- Configure the databases at each site to support the vCenter Server and SRM.

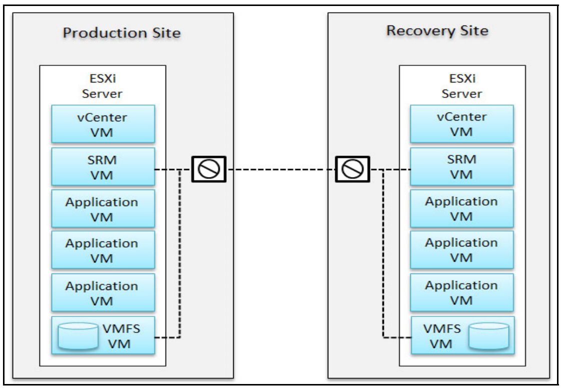

Figure 5. Environment test architecture.

B. *Hardware and Software Configurations*

The hardware and software for the configuration of environment test are listed in Table 1.

TABLE I. THE VMS IN PRODUCTION AND RECOVERY SITE IN KOC TEST ENVIRONMENT.

Hardware	
Production site & Recovery site	VM-1 VM name = "vCenter Server" 2048MB RAM, Hard disk 40 GB VM-2 VM name = " SRM Server" 2048MB RAM, Hard disk 40 GB VM-3 VM name =" DB Server" 2048MB RAM, Hard disk 40 GB VM-4, VM-5, VM-6 VM name ="App1","App2","App3" 2048MB RAM, Hard disk 40 GB
Software	
Production site & Recovery site	VMware: • VMware ESX 4.1 • vCenter Server 4.1 • SRM 4.1 VMs Operating System: • Microsoft Windows Server 2008 • Microsoft Windows XP 64 bit Storage Works: • HP Storage Works P4000 Virtual SAN Appliance SRA: • Version: v1.20.10713 Database: • Microsoft SQL Server 2008

The configuration of both the SRM and vCenter Server requires a database to store information necessary for operation. In order to function properly, both production and recovery sites need a local database. The servers decide the one to operate as the production site and the other to operate as the recovery site.

After SRM has been completely installed on both sites, it is required to connect these sites in order to create a site pair, configure the array managers, configure inventory mappings, and create a protection group and recovery plan at each site. The SRM client plugin is used to administer SRM. Site pairing uses vCenter administrative privileges at both sites.

VI. IMPLEMENTATION OF A DISASTER RECOVERY PLAN

In SRM, the Recovery Plan (RP) consists of certain steps to switch datacenter operation from the production site to the recovery site. The RP ensures that both tests and failovers are executed in a repeatable and reliable manner. Also, it provides approaches to test the BCP and DRP in an isolated environment at the recovery site without impacting the protected VMs at the production site.

A. Testing of Disaster Recovery Plan

After SRM is configured on both the production and recovery sites, the recovery plan can be tested without affecting current services at either site. The test runs recovery plan and, if necessary, configures the two sites for failback so it can restore the services at the production site.

When the test recovery plan is enabled a test network and temporary copy of replicated data at the recovery site are used and the operations are not disrupted at the production site. Testing a recovery plan will complete all the required steps with the exception of the powering down of VMs at the production site and forcing devices at the recovery site to assume mastership of replicated data. A test recovery makes no changes to the production environment at either site.

SRM performs this test with copy-on-write Flashcopies of the mirrored logical drives at the recovery site. The Flashcopy datastores are removed from the recovery host. During the recovery plan configuration, SRM can suspend non-critical VMs on the recovery site to assure me that SRM has enough resources to run the recovery test.

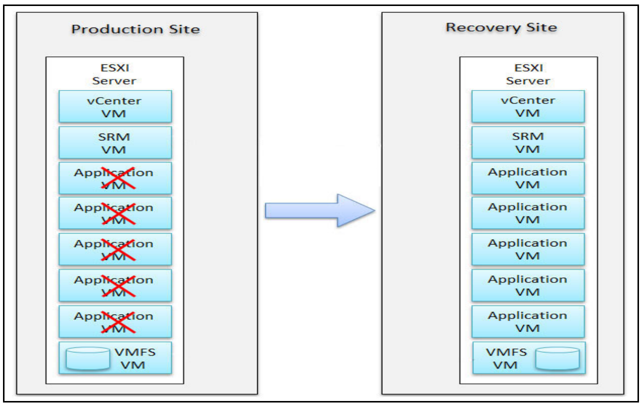

Figure 6. The failover test.

B. Performance Evaluation

Virtualization technology based on a VMware platform conceptually promises potentially powerful, simple and cost-effective solutions to support disaster recovery and business continuity objectives. To verify and validate such assumption when applied to business continuity, the following two proof-of-concept tests were carried out. Herein, the failover test solution proposed two environment tests (Test– 1, Test –2) to estimate and analysis the time for switching the operation from the production site to the recovery site. The two failovers are applied in different VMs environments. Two different policies are used for dynamic virtualized infrastructure operations management as shown in Figure 6.

a) Test-1

This test was conducted by configuring the virtual servers to operate on a sequential schedule basis. In effect, VM1 was scheduled to first start, and VM2 started after the complete loading of VM1. VM2 was scheduled to offer its planned capacity similar to VM3, VM4 and VM5. The test started immediately and sequentially after the previous volume had loaded to its planned capacity. The total elapsed time for the full failover was measured and the results are tabulated in Table II.

TABLE II. TEST-1 ACTUAL TIME.

VM Name	Time **(Min)**	Actual Time **(Min)**
VM1	1.19	1.19
VM2	1.12	2.31
VM3	1.15	3.46
VM4	1.05	4.51
VM5	1.10	5.61
		5.61

It is clear that, VM1 completed loading in 1.19 minutes and VM2 took 1.12 minutes and so on. The total time to load all virtual machines is 5.61 minutes.

b) Test-2

This test was configured in a different way. All virtual servers was scheduled to operate sequentially but almost simultaneously. VM1 was scheduled to start first, and VM2 started after 10 seconds of the loading of VM1. Then, VM2 was scheduled to offer its planned capacity. Similarly each of VM3, VM4 and VM5 was started sequentially 10 seconds after the previous volume has loaded to its planned capacity. The total elapsed time for full failover was measured and the results are shown in Table III.

TABLE III. TEST-2 ACTUAL TIME

VM Name	Time **(Min)**	Actual Time **(Min)**
VM1	1.41	1.41
VM2	1.35	1.45
VM3	1.42	1.62
VM4	1.53	1.83
VM5	2.02	2.42
		2.42

In the above table, VM1 completed loading in 1.41 minutes allowing VM2 to start loading after 10 seconds with duration time of 2.31 minutes. This test was configured to run parallel processing with a 10 second server startup interval for a total elapsed time of 2.42 minutes.

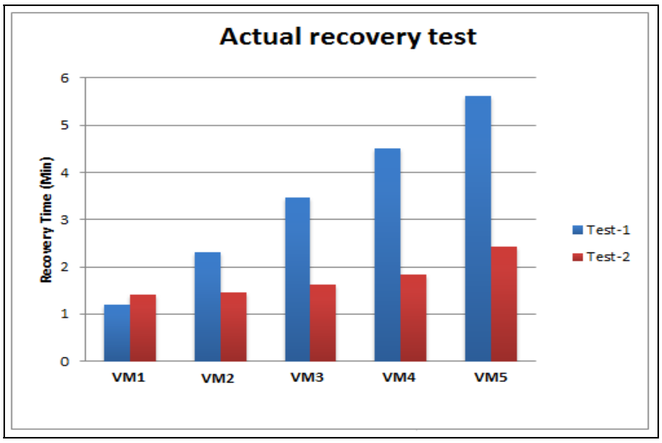

Figure 7. The failover actual times in Test 1 and Test 2.

The results from the two tests confirm the intuitive conclusion that the scheduling arrangements in Test-1 have a more favorable outcome and provide quicker failover and meet more demanding Recovery Time Objectives (RTO) of Service Level Aagreements (SLA) for disaster recovery and business continuity as depicted in Figures 7 and 8. They indicate that the Test-1 scenario is preferable where enterprise application SLAs are stringent and high availability is necessary. The Test-2 is adequate to provide services for scheduled downtime situations because scheduled downtimes are bound by time limit. In Test-2, all virtual servers start at once without wasting time or waiting for other server to start and load fully. Moreover, these tests have validated the premise that VMware virtualization solutions offer proven performance benefits, provide flexibility in operations and offer affordable scalability to data center managers.

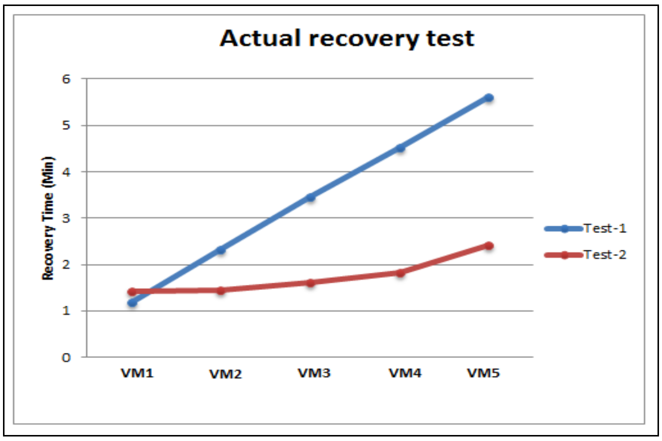

Figure 8. Comparison of failover actual times in Test 1 and Test 2.

VII. CONCLUSIONS

Based on the consideration of the most likely disaster scenarios in this region along with impact analysis and most suitable risk mitigation options, this paper proposed a virtualization solution to sustain IT services at KOC and render these services robust and resilient at all times under all foreseeable conditions. The presented disaster recovery solution consists exploit VMware SRM open-source to reduce the space requirements and maintenance costs at both the production and the recovery sites, where the amount of reduction depends on the server strength. Also, this paper provides a simulation of the testing and monitoring that are performed during a disaster and indicates that the virtualization solution is a cost-effective, simpler and more reliable way to meet business continuity requirements. The solution provides a simplified recovery plan for a virtual environment without any human interactions. Meanwhile, it improves HA and protection of data integrity by synchronous data replication between the production and the recovery sites.

REFERENCES

[1] Anil A. H. Al-Badi and Rafi A. and Ali O. Al-Majeeni and Pam J. Mayhew, "IT disaster recovery: Oman and Cyclone Gonu lessons learned", Information Management & Computer Security, Vol. 17 No. 2, 2009, pp. 114-126

[2] Shaluf, I.M., Ahmadun, F. and Said, A.M. (2003), “A review of disaster and crisis”, Disaster Prevention and Management, Vol. 12 No. 1, pp. 24-32.

[3] McEntire, D.A, “Why vulnerability matters: exploring the merit of an inclusive disaster reduction concept”, Disaster Prevention and Management, Vol. 14 No. 2, 2005, pp. 206-22.

[4] Wainwright, V.L., “Business continuity by design”, Health Management Technology,Vol. 28, 2007, pp. 20-1.

[5] Alexander, D., “Towards the development of a standard in emergency planning”, Disaster Prevention and Management, Vol. 14 No. 2, 2005, pp. 158-75.

[6] Castillo, C., “Disaster preparedness and business continuity planning at Boeing: an integrated model”, Journal of Facilities Management, Vol. 3 No. 1, 2005, pp. 8-26.

[7] Kundu, S.G., “Impact of computer disaster on information management: a study”, Industrial Management & Data Systems, Vol. 104 No. 2, 2005, pp. 136-43.

[8] Jefferson, T.L., “Using the internet to communicate during a crisis”, The Journal of Information and Knowledge Management Systems, Vol. 36 No. 2, 2006, pp. 139-42.

[9] Barbara, M., “Determining the critical success factors of an effective business continuity/disaster recovery program in a post 9/11 world: a multi-method approach”, MSc thesis, Concordia University, Montreal, 2006.

[10] Gorge, M., “Crisis management best practice-where do we start from?”, Computer Fraud & Security, Vol. 6, 2006, pp. 10-13

[11] Atmanand, “Insurance and disaster management: the Indian context”, Disaster Prevention and Management, Vol. 12 No. 4, 2003, pp. 286-304.

[12] T. Wood, E. Cecchet, K. Ramakrishnan, P. Shenoy, J. van der Merwe, and A. Venkataramani, “Disaster recovery as a cloud service: Economic benefits & deployment challenges,” *2nd USENIX Work. Hot Top. Cloud Comput. Boston, MA*, pp. 1–7, 2010.

[13] [2] O. H. Alhazmi and Y. K. Malaiya, “Evaluating disaster recovery plans using the cloud,” *Proc. - Annu. Reliab. Maintainab. Symp.*, 2013.

[14] [3] A. Prazeres and E. Lopes, “Disaster Recovery – A Project Planning Case Study in Portugal,” *Procedia Technol.*, vol. 9, pp. 795–805, 2013.

[15] [4] A. Lenk and S. Tai, “Cloud standby: Disaster recovery of distributed systems in the cloud,” *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, vol. 8745 LNCS, pp. 32–46, 2014.

[16] Shon H., "CISSP All-in-One Exam Guide", Mark Bedell, Fourth Edition, 2007.