World Journal of Advanced Research and Reviews
eISSN: 2581-9615 CODEN (USA): WJARAI
Cross Ref DOI: 10.30574/wjarr
Journal homepage: https://wjarr.com/

(Review Article)

Check for updates

Data mitigation model to establish an emergency recovery solution with an onsite oracle system

Mohammed Sadhik Shaik *

Sr. Software web developer Engineer, Computer science, Germania Insurance, Melissa, Texas.

World Journal of Advanced Research and Reviews, 2021, 10(02), 256-262

Publication history: Received on 12 April 2021; revised on 16 May 2021; accepted on 19 May 2021

Article DOI: https://doi.org/10.30574/wjarr.2021.10.2.0182

Abstract

AWS, or Amazon Web Services, is a cloud computing platform that is adaptable, affordable, and simple to use. The Amazon cloud is a popular location for RDBMSs, or relational database management systems. Oracle Database can be operated on Relational Database Service (Amazon RDS), thus it's important to know how to set it up on AW. With this guide, you will learn all you need to know to run Oracle Database on Amazon RDS, including how to deploy and manage your database, how to handle scalability, performance, backup and recovery, high availability, and security. You will also learn about the benefits of each technique. In this paper, proposed the DM-DATA Model to establish an Emergency Recovery solution with an onsite Oracle system and AWS and to migrate your existing Oracle database to AWS. We provide a strategy for designing an architecture that protects you against hardware failures, datacenter issues, and disasters by using replication technologies stock market data. In the performance analysis, there are several alternatives are choose to optimize the performance of the propose infrastructure with Oracle database based on certain metrics like, disk I/O management, sizing, database replicas, etc.

Keywords: RDS; Amazon; Data Mitigation; Web Services; Data Replication; Oracle

1. Introduction

Oracle Exalogic Elastic Cloud is an all-inclusive platform that combines software and hardware, making it suitable for a wide range of application types and workloads. Deploying large-scale, performance-sensitive mission-critical applications is Oracle Exalogic's strong suit. High levels of isolation across concurrently deployed applications with differing security, reliability, and performance requirements are made possible by combining Oracle Fusion Middleware with Sun hardware [1].

Sun Fire Servers, an Oracle ZFS Storage appliance, plus the necessary InfiniBand and Ethernet networking components make up the Oracle Exalogic Elastic Cloud. Easy management, data services for business continuity, and multi-protocol connection are all features of the Sun ZFS Storage system. Data access protocols supported by the appliance include NFS, iSCSI, and InfiniBand (IB). When it comes to data backup and restoration, the appliance is NDMP compatible [2].An active-passive cluster consisting of two controllers (often called server heads) is an option for configuring the Sun ZFS Storage appliance. Each storage controller in a Sun ZFS Storage appliance is responsible for managing a specific storage pool, which contains all of the storage disks. The storage can be accessed by applications running on the Oracle Exalogic system across the InfiniBand network using the NFSv4 protocol [3].

The Oracle Exadata Database machine is the best solution for hosting the Oracle Database if you need superior database performance. It is both user-friendly and can be deployed rapidly. A "cloud in a box" that includes all the necessary components to host an Oracle Database, the Exadata Database machine is comprised of Database Servers, Oracle Exadata Storage Servers, an InfiniBand fabric for storage networking, and everything else. Database workloads of

* Corresponding author: Mohammed Sadhik Shaik

Copyright © 2021 Author(s) retain the copyright of this article. This article is published under the terms of the Creative Commons Attribution Liscense 4.0.

various types can be efficiently handled by the Oracle Exadata Database machine. This includes data warehousing applications that rely heavily on scans as well as applications that perform online transaction processing (OLTP) at high concurrency. Due to its state-of-the-art hardware components, extensive and intelligent Oracle Database software, and sophisticated Oracle Exadata Storage Server Software, the Oracle Exadata Database machine offers unmatched performance in a dependable and secure environment. Oracle Site Guard, a part of Oracle Enterprise Manager, provides flexible and seamless orchestration of failovers and switchovers across disaster recovery sites, ensuring that enterprise installations see little downtime.

Oracle Site Guard's disaster recovery automation features make failovers and switchovers automated, meaning no human interaction is required and no human error can occur. One of Oracle Site Guard's many strengths is the ease with which it interfaces with other platforms, such as Exalogic and Exadata [4]. Although there are other conceivable disaster recovery topologies, Oracle suggests a topology that is comparable to the one described in this paper. That is, both locations should include a pair of Exalogic and Exadata machines, although their configuration and capacity may differ. Easy management, maintenance, and automated disaster recovery procedures are key components of the topology outlined in this paper's architecture to sustain SLAs.

2. Litertaure review

The increased frequency of natural disasters such as earthquakes, floods, pandemics, and droughts is directly attributable to the quickening pace of climate change. In addition to posing a serious danger to infrastructure, agriculture, and the financial system, these severe weather disasters endanger the lives of millions of people. Inadequate or delayed information transmission, allocation of suitable resources, and risk minimization are the primary causes of casualties in disaster management. In order to prevent these kinds of losses, it is crucial to put technology in place that can aid in the efficient management of various parts and stages of emergency, disaster, and extreme weather situations [5].

The cloud makes it possible for many organizations to share storage space, and artificial intelligence (AI) improves the cloud's ability to glean useful information before a disaster strikes so that alerts can be developed, and during the disaster itself so that rescue operations can be carried out and priorities can be set [6].

Cloud systems powered by artificial intelligence help communities recover more quickly after disasters, allowing businesses to go back to work as soon as possible and be ready for anything the future throws at them. Most public systems have a hard time keeping individuals informed during times of harsh weather and natural disasters regarding available emergency shelters, casualties, government laws, medications, and diagnostic and treatment choices. Furthermore, disinformation has the potential to sow discord and panic among the populace, which in turn can lead to a scarcity of goods as a consequence of frantic purchasing, price spikes, and prejudice [7].

Discrimination manifests itself in a variety of ways, including unequal access to treatment, forced relocation, and aid distribution. For this reason, it is of the utmost importance that governments and agencies broadcast transparently accurate and verified information that vulnerable populations may use during times of extreme weather. As an example, a location-and mobile-driven system helped local authorities evacuate millions of people during the recent floods in India [8]. When it comes to disaster relief, emergency operations, and severe weather, a lot of different groups work together. This includes government agencies like fire departments and hospitals as well as NGOs like Oxfam and the Red Cross. It also includes research and development centers and businesses like the oil and gas industry, mining, construction, and fishing.

An intuitive, real-time data-gathering technology that facilitates frictionless information sharing is essential for effective collaboration [9]. Mismanagement or a disjointed approach to handling disasters and emergencies is commonplace since each stakeholder typically employs their own data and method to crisis management. Combining the information supplied by all stakeholders and extracting useful data for further action using an AI and cloud-based system can lead to a faster execution and recovery time, which in turn can save more lives.

There are a number of ways in which cloud computing and artificial intelligence (AI) can help in disaster preparedness and response. On the one hand, AI can make it easier for public and private organizations, as well as victims and volunteers, to work together before, during, and after an event [10].

In this approach, catastrophe, severe weather, and emergency operations can be better managed with the use of artificial intelligence and collaborative platforms hosted in the cloud. By facilitating the accurate and rapid processing of information among stakeholders and the alignment of inter-organizational operations, AI and cloud-based collaborative platforms can enhance disaster response. Because of this, damage assessments will take less time, and supplies and aid can be delivered more quickly during times of extreme weather [11]. Our primary concern in times of crisis is always finding a way to escape the danger as soon as possible, whether it's through evacuating to a safer location (during floods or storms, for example) or obtaining medical treatment to ward off an illness [12].

Along these lines, many businesses are incorporating green production practices into their commodities distribution strategies and are monitoring their progress with the help of emerging technologies like AI. By improving materials and using robots to cut down on waste and energy consumption, artificial intelligence (AI) is also used to lessen the negative effects of production on the environment and to lessen the carbon footprint. It is also possible to understand the spread of meteorological phenomena and plan for their mitigation, readiness, reaction, and recovery using cloud-based collaborative platforms and AI tools [13]. Climate services, which can detect extreme events, remote sensing, which may offer satellite images and aid in damage assessments, and international charters are all examples of such tools. These solutions can significantly enhance situational awareness among stakeholders, distribution of resources, timeliness of damage assessments, and tracking of notifications during evacuations.

One crucial application of this technology could be to automatically notify hospitals when patients require specific treatment [14]. In locations where people are often unable to leave their homes due to weather or other emergencies, facial recognition technology can help track their whereabouts, while unmanned aerial vehicles (UAVs) can transport groceries and other household goods. Artificial intelligence (AI) and cloud-based collaboration tools can also detect rulebreakers in severe weather scenarios and periodically notify them.

From a business point of view, there are two things that are absolutely necessary in the event of an emergency or severe weather: first, the ability to store, secure, and make data available to drive choices; and second, the power of the organization to analyze and process information. In this case, AI need a big dataset in order to provide reliable recommendations and forecasts; combining AI with a cloud-based collaboration platform that can translate between spoken and written language is the best way to do this. Artificial intelligence enhances the accuracy of weather predictions through the use of pattern recognition, natural language processing, and forecasting techniques, while cloud computing provides a platform for assistance in the event of emergencies and extreme weather [15].

In the event of an emergency or severe weather, mobile applications provide convenient access, simplicity of use, and convenience of update. In the modern era of widespread smartphone ownership and widespread internet access, a mobile application—a platform that is highly linked, secure, real-time, flexible, and agile—has the capacity to aid in the event of an emergency or severe weather event. When dealing with emergency situations or severe weather, it is important to have the ability to synchronize data continuously so that it can be quickly shared across stakeholder networks through cloud platforms.

This data can be in various forms, including text, video, audio, numbers, and photos. Intelligent technologies, like AI, can analyze this data. Technologies that mitigate issues before, during, or after a disaster have been the focus of previous research [16]. But research hasn't looked at how to use intelligent technology in an integrated fashion to handle all the stages of a disaster cycle. Research also fails to reevaluate the current state of affairs in light of technological advancements in the event of impending disasters. The following research issue is investigated in this work to fill these gaps: When it comes to responding to emergencies, severe weather, and disaster relief efforts, how might artificial intelligence and collaboration platforms in the cloud be of assistance? Researchers create a semi-structured interview schedule to collect data from people who have used artificial intelligence and cloud-based collaboration platforms in times of crisis or severe weather. Therefore, OIPT is appropriate for this research since it is crucial to have systems and technology ready to deal with uncertainties that arise as a result of a disaster [17]. Moreover, following data collection, a thorough coding procedure was employed to produce themes, propositions, and a unique framework (4-AIDE) supported by evidence.

3. Disaster recovery for oracle supercluster

The next sections provide an overview of Oracle SuperCluster and the proposed disaster recovery strategy to guarantee application availability to the highest possible level.

3.1. Oracle SuperCluster Overview

Oracle SuperCluster is a multipurpose designed system that can run mission-critical enterprise applications and swiftly build cloud services with great efficiency, cost savings, and performance. With Oracle SuperCluster, you get an infrastructure that has been through extensive testing and is prepared for deployment. It includes SPARC servers, Exadata storage servers, the ZFS storage appliance, InfiniBand technology, and Oracle Solaris. Multitier business applications—which consist of web, database, and application components—are well-suited to this technology. The complete Oracle software solution stack, as well as third-party applications and customer-developed software, can all be hosted within a single rack enclosure by Oracle SuperCluster. Oracle SuperCluster combines scalable and highly available technologies with industry-standard hardware, with end-to-end high availability as its basic design. Includes Oracle Real Application Clusters (RAC) and Oracle Clusterware, Oracle Database 11g enables high availability and database failover. Programs are guaranteed to be highly available using Oracle Solaris Cluster. Some of the most important pieces of hardware are built to withstand failures; these include SPARC servers, Oracle ZFS Storage Appliance, and Oracle Exadata Storage Servers. Oracle Solaris ZFS file system capabilities like self-healing, triple-RAID parity, and triplemirroring contribute to data availability, together with data protection features like dual controllers and ECC memory. Databases and applications are the two basic kinds of domains that SPARC servers can accommodate. Oracle Database 11g Release 2 and later, in conjunction with Oracle Exadata Storage Servers, are the only authorized users of Database Domains.

3.2. Disaster recovery strategy for oracle supercluster

Even though Oracle SuperCluster is designed for high availability, enterprise deployments still need to be protected from unexpected disasters and natural calamities. Oracle SuperCluster uses a mix of technologies to offer disaster recovery assistance for databases and applications running on this platform (see Figure 1). Oracle GoldenGate and Oracle Active Data Guard are our top recommendations for replicating database material. We advise using ZFS replication for apps and unstructured data. We also suggest Oracle Enterprise Manager and Oracle Solaris Cluster Geographic Edition for controlling the complete disaster recovery system.

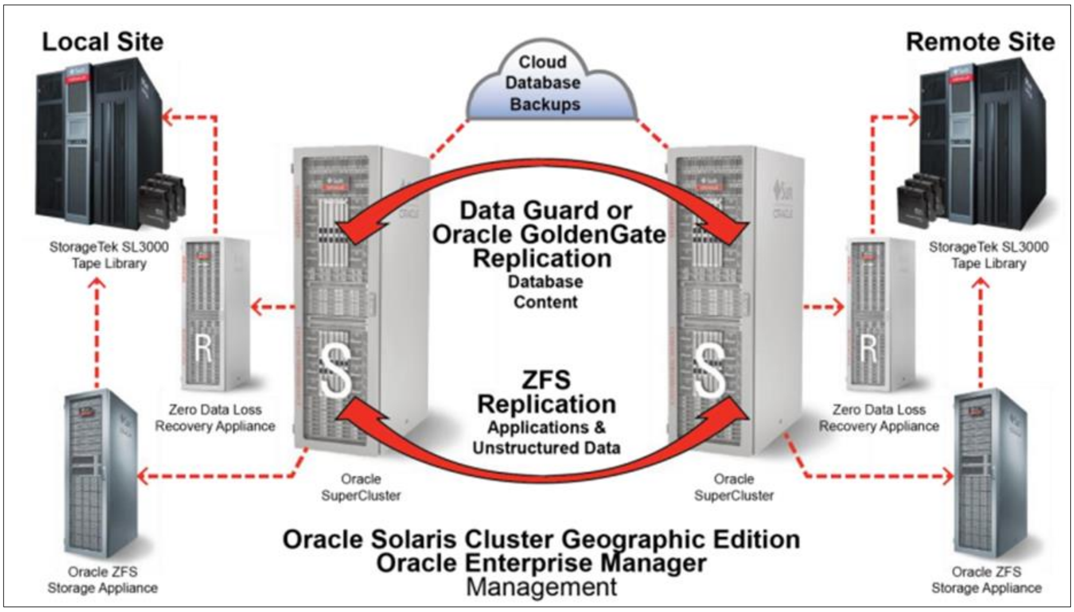

Figure 1 One viable option for a disaster recovery architecture is to include Oracle SuperCluster

On the Oracle SuperCluster platform, the shared file systems of the Oracle ZFS Storage Appliance hold applications and non-database data, commonly referred to as unstructured data.

Applications such as Oracle enterprise apps, apps developed by third parties, and apps built specifically for use with the Oracle SuperCluster platform can all be part of this data. This data's disaster recovery plans make use of Oracle ZFS Storage Appliance's remote replication capabilities. When compared to older offline backup designs, the recovery time after a disaster is much reduced when a copy of the main data is kept at a remote location. Oracle SuperCluster's Oracle ZFS Storage Appliance has a dedicated 1 Gigabit Ethernet (GbE) connection for replication, therefore no more gear is needed for that. The Oracle ZFS Storage Appliance, an internal component of Oracle SuperCluster, comes with replication and cloning licenses, which are sold separately. Every one of the remote and local sites' external Oracle ZFS Storage Appliances, however, need its own license.

All databases, both Oracle and third-party, that are compatible with Oracle Solaris 11 and Oracle's SPARC servers are available in Oracle SuperCluster. For database storage, Oracle Exadata Storage Servers are accessible to instances of

Oracle Database 11g Release 2 (or later) running in the Database Domain of Oracle SuperCluster. Oracle Exadata Storage Servers are inaccessible to databases developed for or used by non-Oracle platforms or previous versions of Oracle Database that operate in the Application Domain of Oracle SuperCluster. The data source for these databases can be an external Fibre Channel SAN or an Oracle ZFS Storage Appliance.

Data Guard and Oracle Active Data Guard are the most popular recommendations for Oracle Database disaster recovery. Oracle GoldenGate is the best option for non-Oracle database environments that require disaster recovery, for replication across various Oracle Database releases, or for bidirectional replication where both replicas must be open in read-write mode simultaneously.

You can utilize ZFS replication with databases that use Oracle ZFS Storage Appliance as well. Even though they aren't ideal, non-Oracle database replication methods are nonetheless available for outdated setups. Oracle SuperCluster's 10 GbE ports are usually shared across users, applications, and database replication. Nevertheless, individual ports can be set up to accommodate performance or other company needs.

Extra hardware for the wide area network (WAN) may be required to ensure company continuity in the case of total site failures. The standby site must be forwarded to users in order to keep availability. Users can be automatically or manually sent to the application tier at the standby site through a Domain Name System (DNS) failover using a WAN traffic management. Similarly, a database failover can transfer the standby database to the main production role. For details on how to automate a full site failover, refer to the Oracle Database high availability best practices manual.

4. Disaster recovery architecture

This project's reference architecture was based on the topology given in the Oracle Business Intelligence Enterprise Deployment Guide for Oracle Fusion Middleware. This reference design was adjusted for Exalogic and Exadata machine deployment in accordance with the guidelines in the Oracle Fusion Middleware Exalogic Enterprise Deployment Guide. Oracle Fusion Middleware Disaster Recovery Guide and Oracle Enterprise Manager Cloud Control Lifecycle Management Administrator's Guide were followed to modify it for Oracle Site Guard-based disaster recovery. Following best practices and based on known Oracle high-availability and security solutions, an Enterprise deployment is the recommended configuration for Oracle Exalogic. It is intended to run large-scale, mission-critical corporate software applications.

4.1. Overview of Topology

A high-level view of the topology used in the deployment and testing of this paper is shown in figure 2 below.

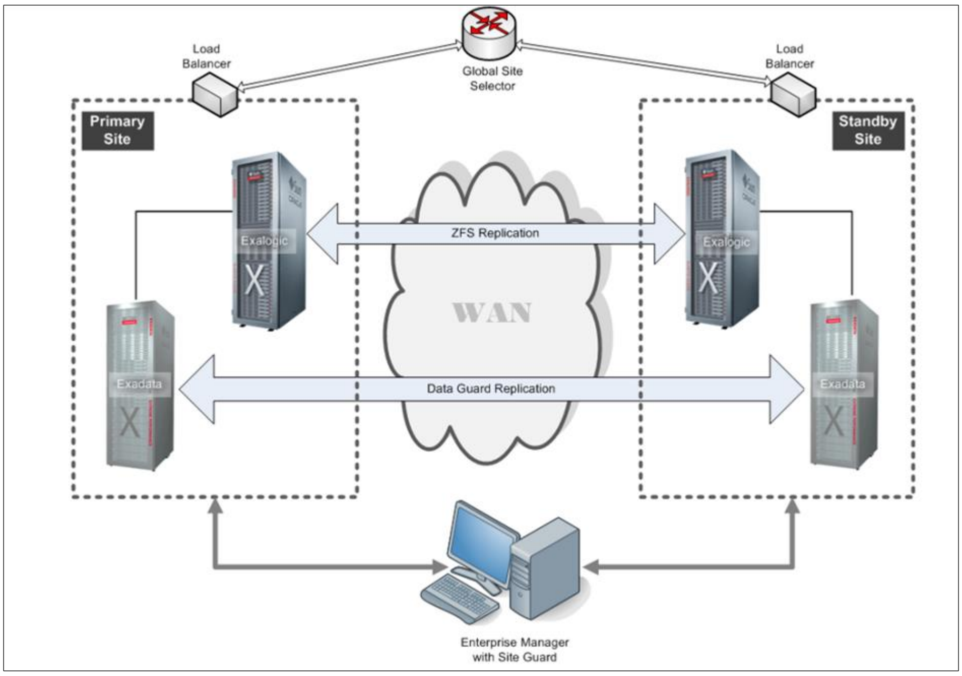

Figure 2 Overview of the topology

Both the main and backup locations include:

- Two Exalogic-based Oracle HTTP Servers. (domain1, domain2)
- One Oracle Exalogic-based Oracle Weblogic Administration Server.
- The application host Oracle Exalogic-managed components and two Oracle Weblogic Managed Servers are part of the third system. (application hosts 1 and 2).
- Application data is stored in a single-node Oracle RAC database that is run on an Oracle Exadata Database machine. (client database hosts1, client database hosts2)
- It was on the Exalogic machine's shared storage that the Oracle HTTP Server and Oracle Weblogic Server configuration files and binaries were installed.

Along with the previously mentioned site-specific hosts, one standalone server called emcchost runs Oracle Enterprise Manager Cloud Control with Oracle Site Guard plug-in. This server is utilized for managing both sites. While all of the hosts—webhosts and apphosts—used in the deployment described in this paper were virtual machines (vServers), they could also be configured on physical servers within the Exalogic system.

Below is a schematic illustrating the host deployment process at each site.

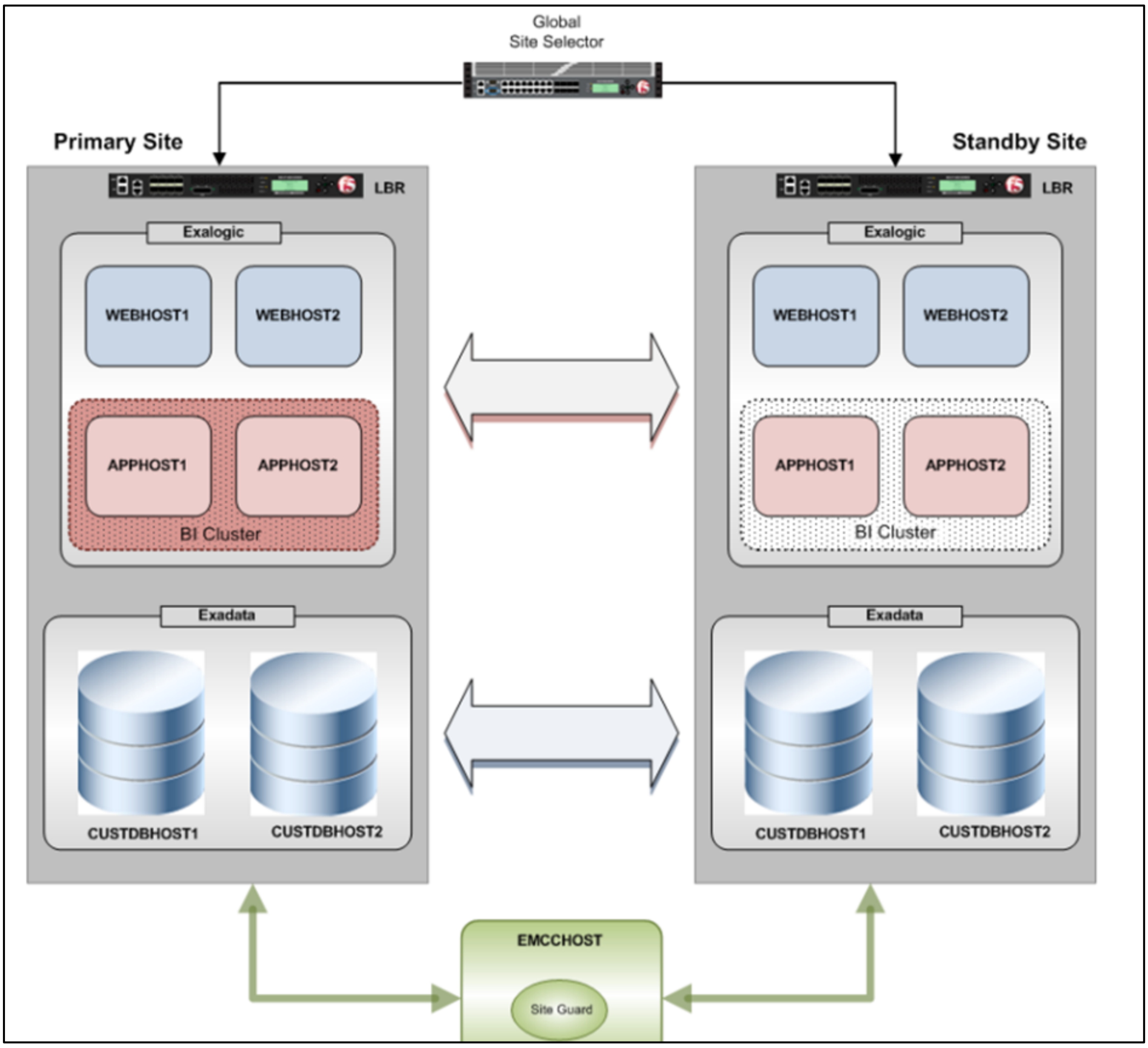

Figure 3 How the hosts are deployed at each site

5. Conclusion

A standby site, located in a separate geographic region from the production site, is typically put up as a disaster recovery solution. Periodically or continuously, all data is replicated to the standby siteEverything from application data and configuration files to metadata and database records falls under this category. In the event of an emergency, the backup location can take over operations.

Oracle-Enhanced Disaster Recovery Solution offers next-gen data security using components from Oracle's comprehensive software and hardware technology stack, which includes Data Guard, Oracle GoldenGate, Oracle Zero Data Loss Recovery Appliance, and Oracle ZFS Storage Appliance. This solution safeguards the cluster's middle-tier applications and components by utilizing the replication capability of Oracle ZFS Storage Appliance. For databases included in Oracle SuperCluster deployments, disaster recovery is provided via either Data Guard or Oracle GoldenGate. Management of the complete disaster recovery solution is provided by Oracle Solaris Cluster Geographic Edition and Oracle Enterprise Manager. Integration of non-Oracle databases1 and applications, as well as old databases, is also possible with the help of third-party replication tools.

References

[1] Ahmad, M.I., Ma, H., 2020. An investigation of the targeting and allocation of post-flood disaster aid for rehabilitation in Punjab, Pakistan. Int. J. Disaster Risk Reduc. https://doi.org/10.1016/j.ijdrr.2019.101402.

[2] Ahn, J., Son, H., Chung, A.D., 2021. Understanding public engagement on Twitter using topic modelling: the 2019 Ridgecrest earthquake case. International Journal of Information Management Data Insights. https://doi.org/10.1016/j.jjimei.2021.100033.

[3] Aitsi-Selmi, A., Murray, V., Wannous, C., Dickinson, C., Johnston, D., Kawasaki, A., Stevance, A.S., Yeung, T., 2016. Reflections on a science and technology agenda for 21st century disaster risk reduction. International Journal of Disaster Risk Science 7 (1), 1–29.

[4] Akter, S., Wamba, S.F., 2019. Big data and disaster management: a systematic review and agenda for future research. Ann. Oper. Res. 283 (1–2), 939–959.

[5] Al Qundus, J., Dabbour, K., Gupta, S., Meissonier, R., Paschke, A., 2020. Wireless sensor network for AI-based flood disaster detection. Ann. Oper. Res. https://doi.org/10.1007/s10479-020-03754-x.

[6] Aldunce, P., Beilin, R., Handmer, J., Howden, M., 2014. Framing disaster resilience. Disaster Prev. Manag. 23 (3), 252–270.

[7] Alford, J., 2020. Respond, recover and reimagine: A framework for responding to the coronavirus and beyond. Assessed on 22 June 2021. https://www.sas.com/en_ae/insights/articles/analytics/respond-recover-and-reimagine-a-framework.html.

[8] AlJahdali, H., Albatli, A., Garraghan, P., Townend, P., Lau, L., Xu, J., 2014. Multitenancy in cloud computing. In: IEEE 8th International Symposium on Service Oriented System Engineering. IEEE, pp. 344–351.

[9] Ansari, H.R., 2014. Use seismic colored inversion and power law committee machine based on imperial competitive algorithm for improving porosity prediction in a heterogeneous reservoir. J. Appl. Geophys. 108, 61–68.

[10] Baum, A., Fleming, R., Davidson, L.M., 1983. Natural disaster and technological catastrophe. Environ. Behav. 15 (3), 333–354.

[11] Baxter, P., Jack, S., 2008. Qualitative case study methodology: study design and implementation for novice researchers. Qual. Rep. 13 (4), 544–559.

[12] Behl, A., Dutta, P., 2019. Humanitarian supply chain management: a thematic literature review and future directions of research. Ann. Oper. Res. 283 (1), 1001–1044.

[13] Behl, A., Dutta, P., Luo, Z., Sheorey, P., 2021. Enabling artificial intelligence on a donation-based crowdfunding platform: a theoretical approach. Ann. Oper. Res. https://doi.org/10.1007/s10479-020-03906-z.

[14] Bharosa, N., Lee, J., Janssen, M., 2010. Challenges and obstacles in sharing and coordinating information during multi-agency disaster response: propositions from field exercises. Inf. Syst. Front 12 (1), 49–65.

[15] Blaikie, P., Cannon, T., Davis, I., Wisner, B., 2014. At Risk: Natural Hazards, People's Vulnerability and Disasters. Routledge, London.

[16] Bradbury, D., 2011. Data mining with LinkedIn. Comput. Fraud Secur. 2011 (10), 5–8.

[17] Bui, X.N., Nguyen, H., Le, H.A., Bui, H.B., Do, N.H., 2020. Predictions of blast-induced air over-pressure in open-pit mine: assessment of different artificial intelligence techniques. Nat. Resour. Res. 29 (2), 571–591.