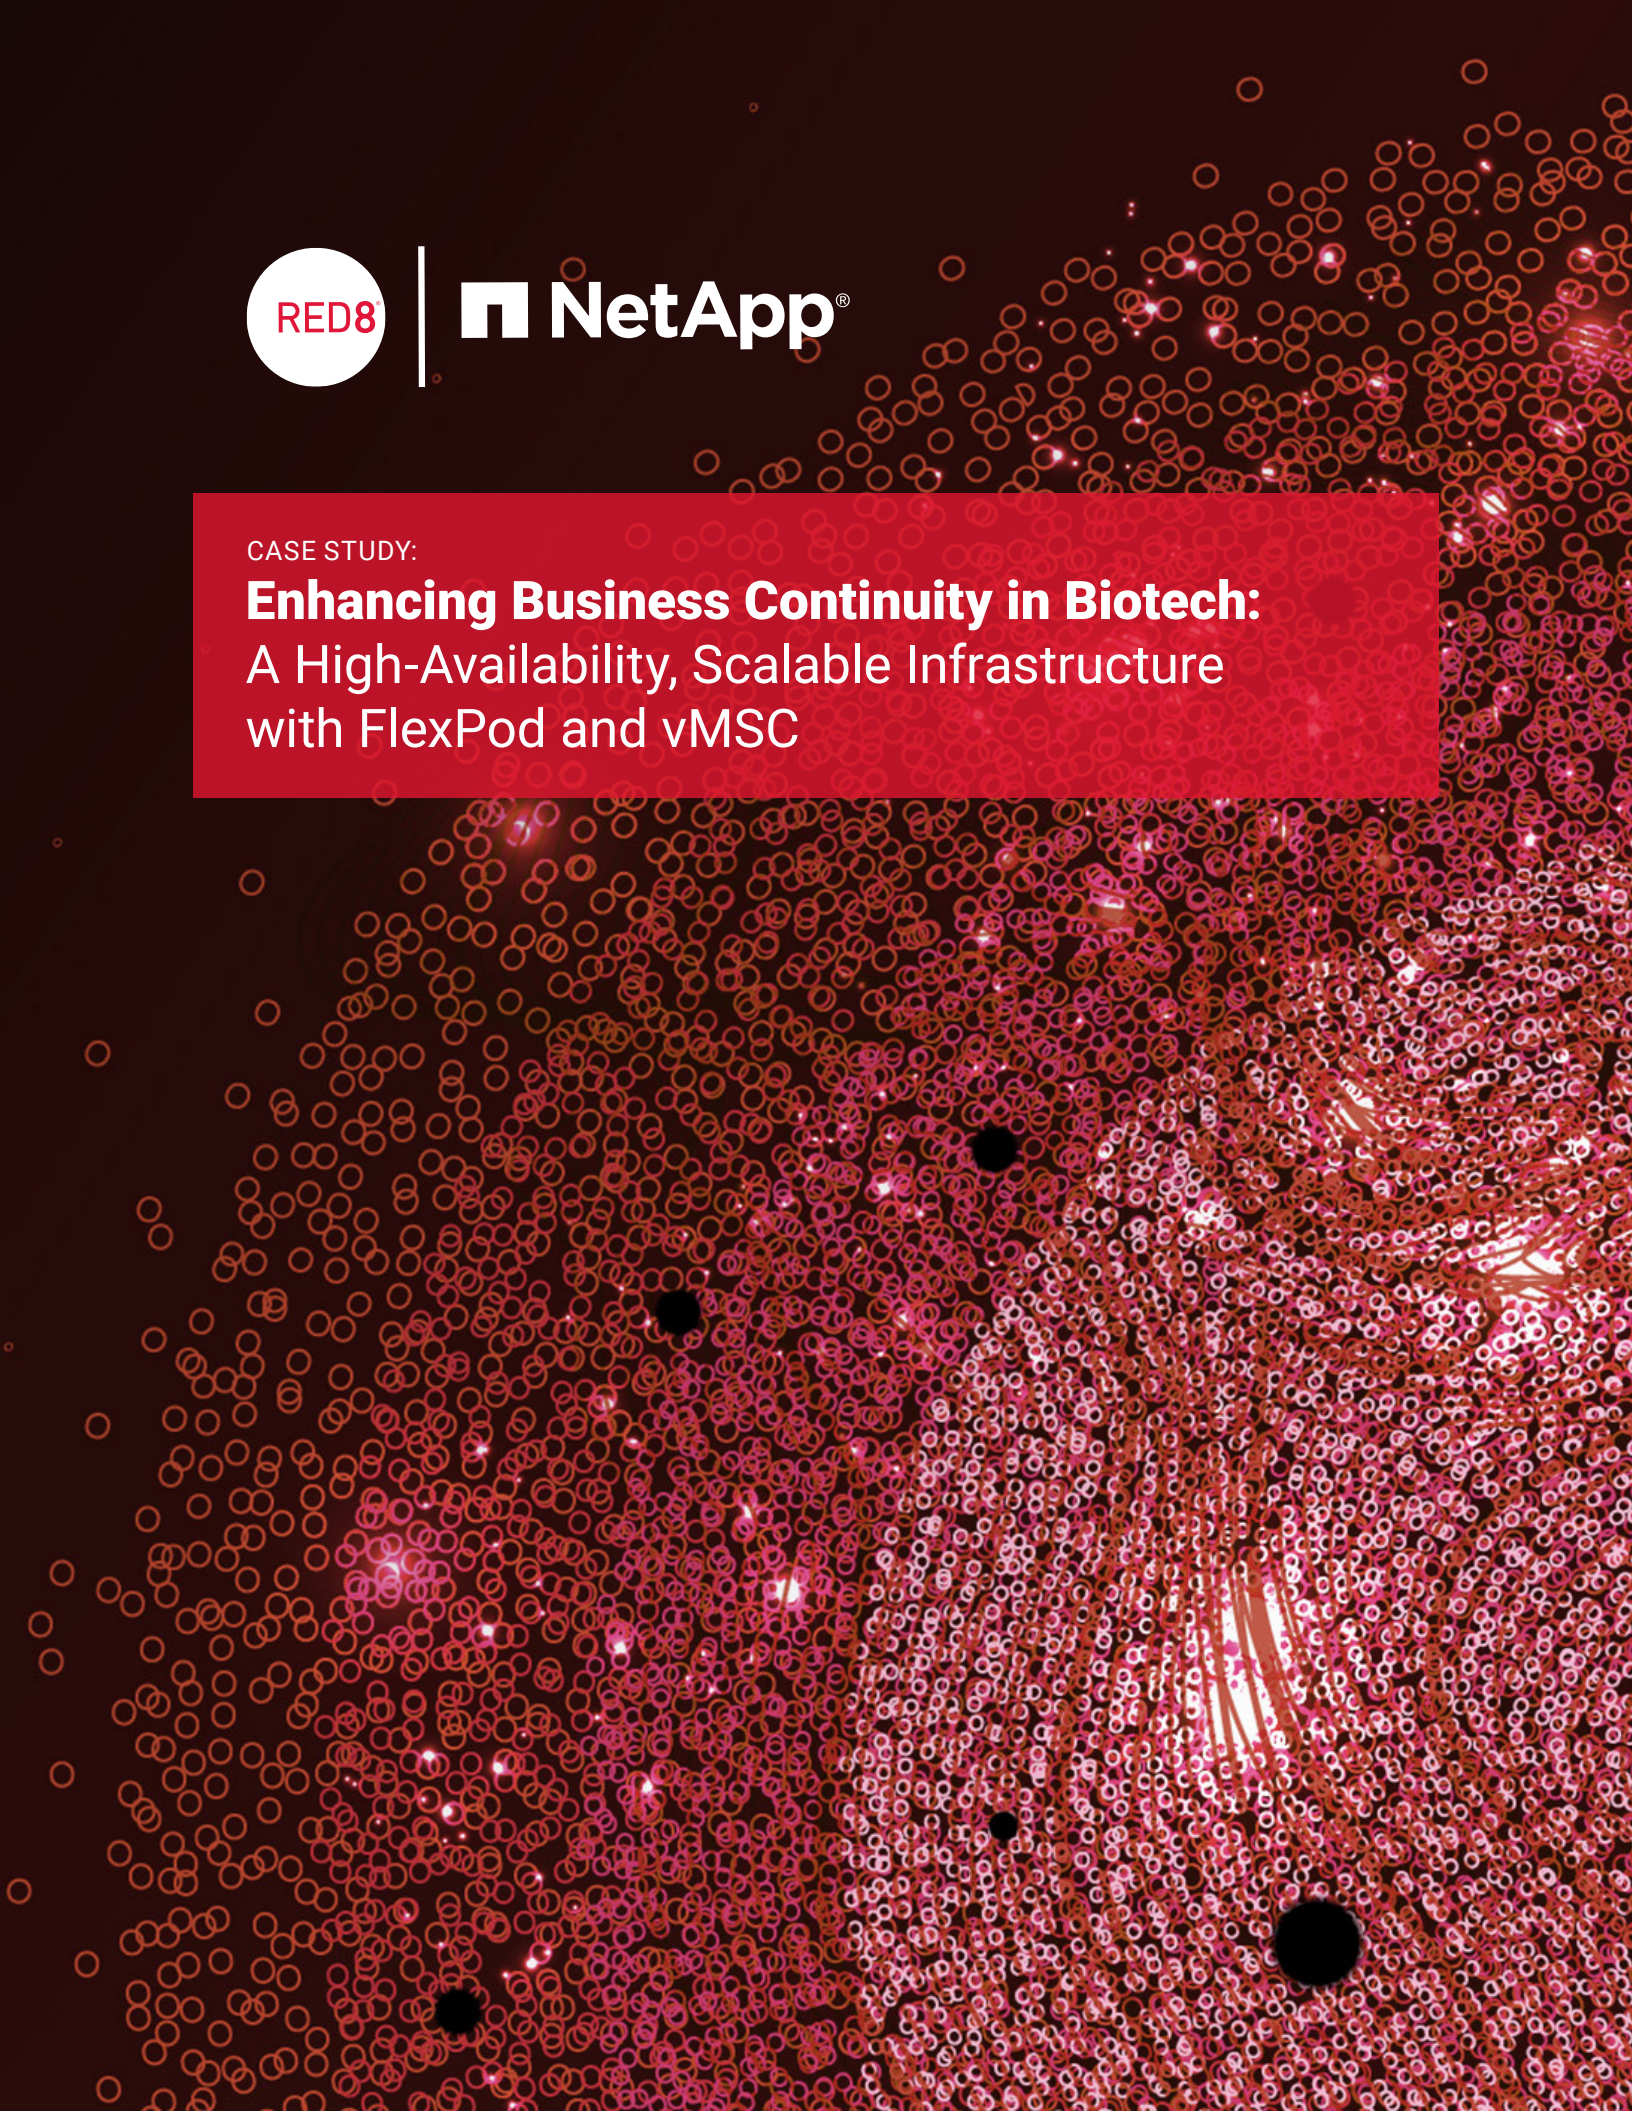

CASE STUDY:

Enhancing Business Continuity in Biotech: A High-Availability, Scalable Infrastructure with FlexPod and vMSC

Customer Overview

The customer is a leader in the biotechnology sector that focuses on developing medicines that improve people's quality of life. The firm is led by scientists who can consistently productize their research into FDA-approved medicines leveraging a significant genetic sequencing effort. Performing this research requires robust information technology systems, and the importance of time-to-market drives high regard for business continuity and data availability. In fact,business resiliency and data access improvement are strategic initiatives.

Problem Statement

Recently the customer's business has been in high-growth mode - developing several different products, expanding research, and manufacturing facilities. The business leaders decided to take advantage of the additional facilities to enhance their business continuity capabilities and eliminate possible disruptions while expanding their computing and data management capabilities. The importance of system availability and continuity drove the business leaders to invest into future business technology.

Before this growth, the customer ran hundreds of servers supporting numerous Microsoft and Linux applications, including SAS, custom software applications, large data marts and data warehouses, and manufacturing and distribution applications. All these applications had been operating seamlessly on a NetApp FlexPod environment initially designed and deployed by Red8 just over a year ago.

The originally deployed infrastructure included 12 Cisco UCS blades on two chassis, Cisco switches, a NetApp AFF-A300 with shelves of solid-state drives, and an end-to-end 40 Gigabit Cisco network infrastructure (see **Figure 1**). The computing platform ran VMware 6.7 Update 3 over NFS with iSCSI SAN boot so that server blades could be swapped out easily without needing to deal with onboard blade media. NFS was used as the virtual machine (server) file system for the operational efficiencies of managing NFS as opposed to VMFS and FC or iSCSI.

Figure 1. *Original FlexPod Infrastructure environment*

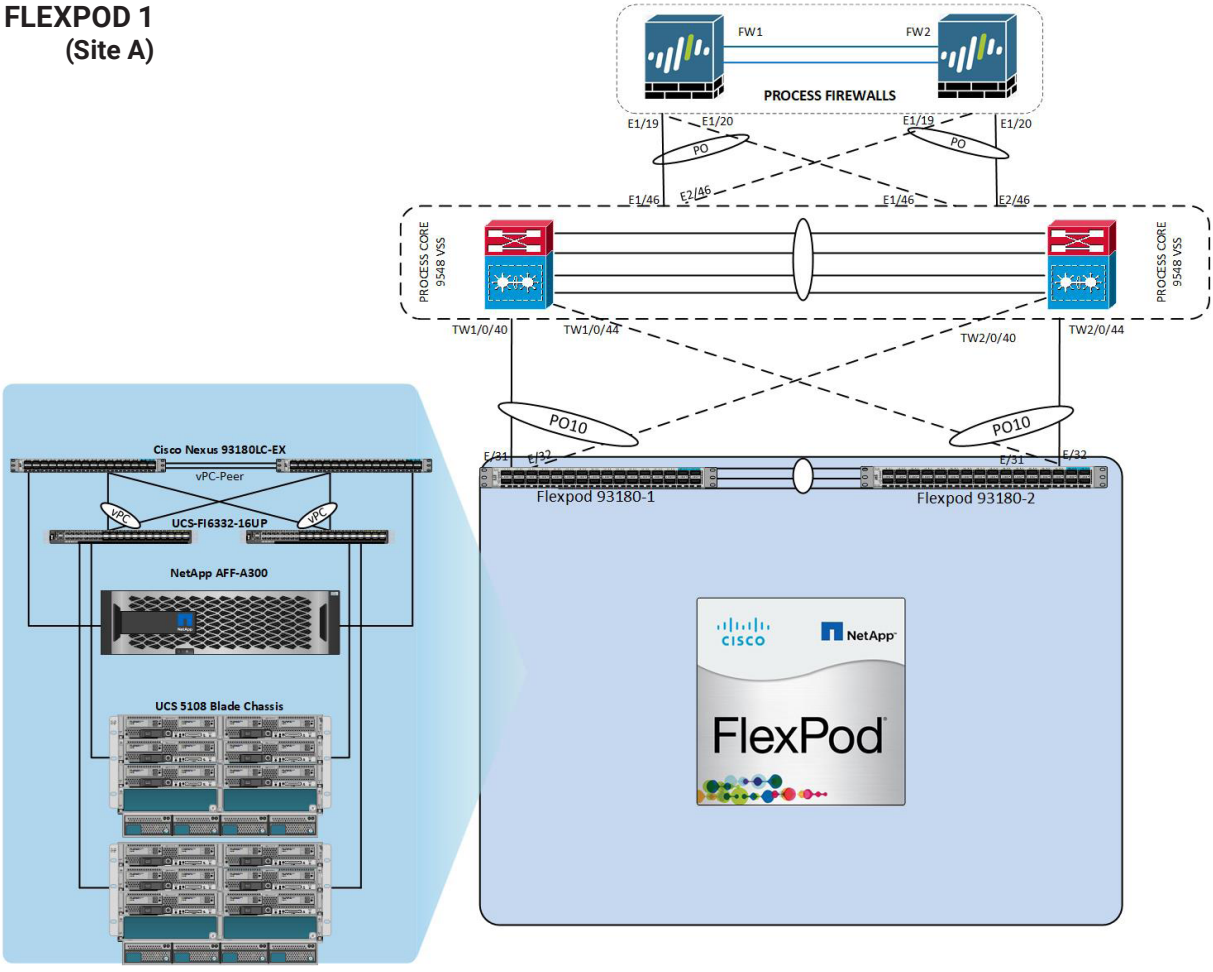

Solutions Considered

The business was experiencing high growth, and the emphasis on business continuity and data availability drove the customer to advance their failover capabilities more. To this end, the customer envisioned a solution that would stretch the existing data center facility to a new one they were building approximately a mile away.

Given the existing infrastructure, the customer considered the following potential solutions:

1. Separate FlexPod and use vSphere Site Recovery Manager (SRM) and SnapMirror for failover.
2. Separate FlexPod and use application dependencies to run applications (i.e., Microsoft Geo-Cluster, vSphere HA)
3. vSphere Metro Storage Cluster (vMSC) with NetApp MetroCluster™

Table 1 describes some of the pros and cons of the above options.

Table 1. *Pros and cons of the various options considered.*

Option	Pros	Cons
Separate FlexPod and use SRM and SnapMirror for failover	All done through vCenter; no underlying storage requirements. No latency requirements or distance limitations.	More complex because one has to set up failover rules (e.g., one must do a lot of networking configuration, boot order, etc.). SRM licensing is required
Separate FlexPod and use application failover to run applications (i.e., Microsoft Geo-Cluster, vSphere HA)	No additional licensing. Can customize the failover of and DR test each application	Not everything would be highly available (multi-site HA). Certain VMs wouldn't be set up for HA; they'd require a manual stand up. Applications are running in both locations (active/passive) Large VMs over eight vCPUs cannot use vSphere HA
vSphere Metro Storage Cluster (vMSC) with NetApp MetroCluster	Low complexity Easy to move workloads between locations. All 24 blade servers could be in the same vSphere cluster.	Non-redundant applications will see a short interruption if a VM must be restarted. Latency requirement of 10ms (and max distance). Need extra switches.

Solution Selected

Ultimately, the customer selected the third option (vMSC) because it provided the highest availability with the least complexity. To accomplish this required the exact same hardware as originally configured for Site A. Once the hardware was in place, the entire environment was reconfigured as a MetroCluster to support vMSC. From the application/VM standpoint, vSphere's recoverability features determined the hosts to which each VM would be assigned (vSphere DRS and HA). The fact that all 24 servers could be included in the same vSphere cluster made it very easy to administer and allowed VMware vSphere Distributed Resource Scheduler (DRS) to take over the management of the location of the workloads.

A VMware vMSC configuration is a specific storage configuration that combines replication with array-based clustering and is typically deployed in environments where the distance between data centers is limited. Since the new data center location is not far from the current data center, vMSC provided an optimal solution. In addition, vMSC infrastructures offer the same benefits that vSphere HA clusters provide to a local site but in a geographically dispersed model with two data centers in different locations.

A vMSC infrastructure is essentially a stretched cluster. The architecture is built on the premise of extending what is defined as "local" in terms of network, storage, and computing to enable these subsystems to span geographies, presenting a single and common base infrastructure set of resources to the vSphere cluster at both sites. It stretches storage, network, and computing between sites.

Figure 2. *Design of final solution.*

VMware vMSC w/NetApp MetroCluster

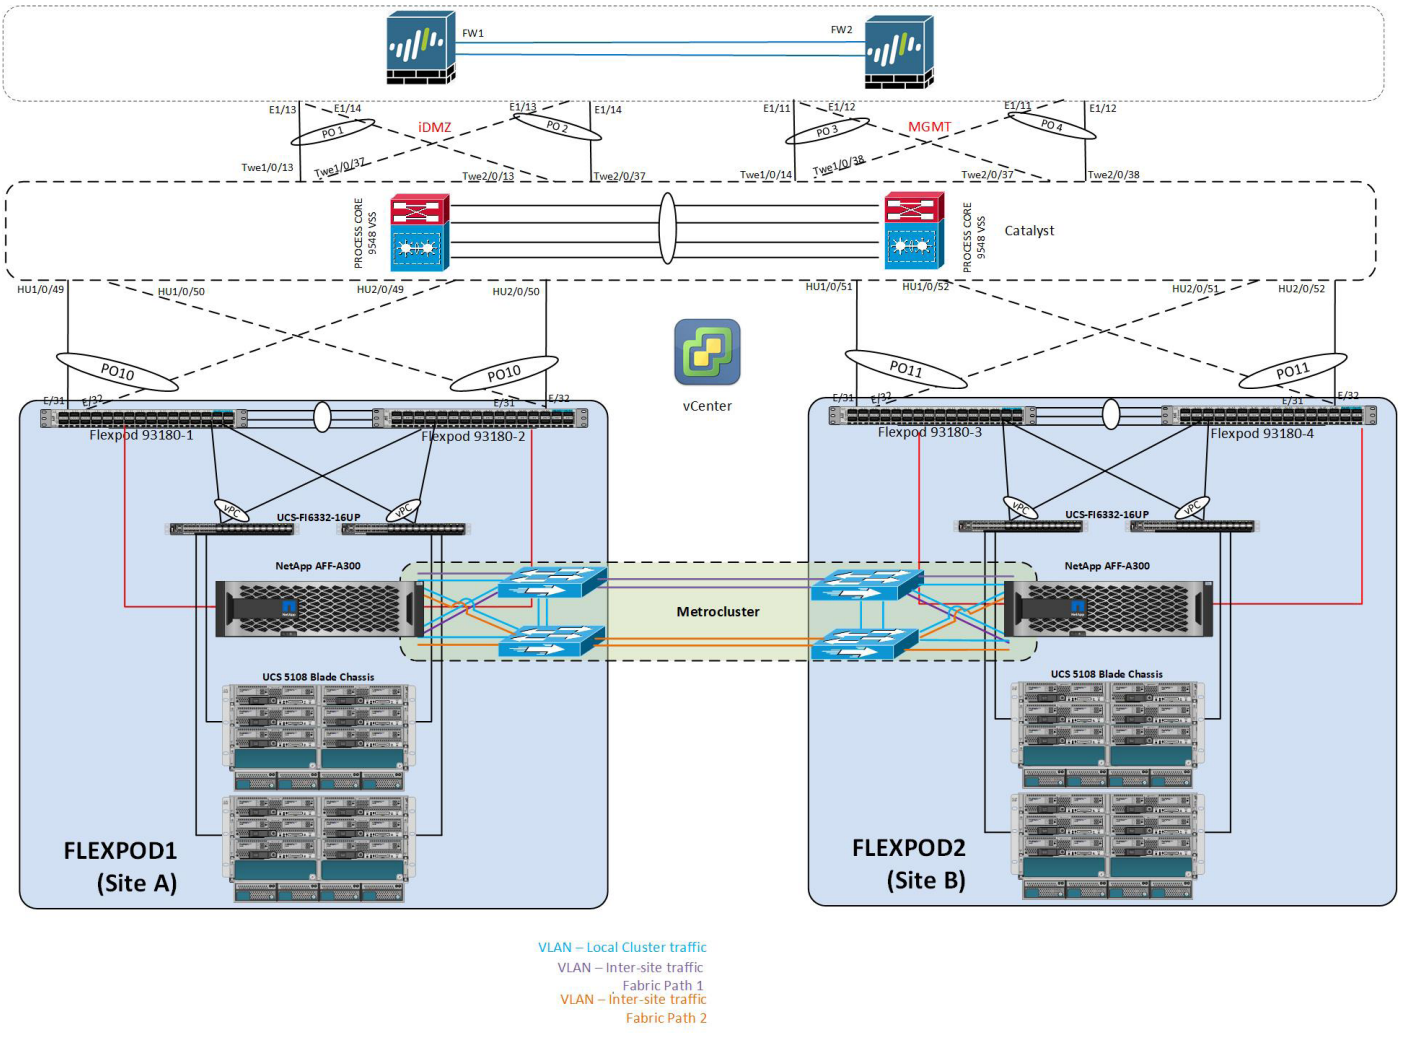

Results

Since the new data center is still under construction, the new solution was built in the current data center. Once the new data center is online, the customer will move a portion of the environment to the new data center while maintaining 100% uptime, and the workloads would automatically be dispersed between both MetroCluster "sites."

Red8 and the customer performed the following acceptance tests for the initial deployment.

1. Manual "switch over" from old site to new site platform with applications.
    a. The failover happened within seconds leaving the customer wondering whether the switchover test had occurred at all.
    b. Also performed a switch back with the same results.
2. Simulated power failure by hard powering off both storage nodes in one site. Again, the switch over to the other "site" happened automatically thanks to an independent tiebreaker (a part of the solution that checks for the health of the environment; the Tiebreaker is a Linux virtual machine that monitors the environment and decides which side will take over during a failure).
3. Storage operations system "ONTAP" upgrade using NetApp's Automated Non-Disruptive Upgrade. This also completed with no disruption to the applications.

Given the success of this solution, the customer is considering additional MetroCluster deployments, including the conversion of FlexPods previously deployed in a different configuration. The customer feels completely comfortable with the imminent move of the equipment to the new site and the added resiliency off their applications is a welcome bonus.

Summary of Results

Currently, the customer is extremely impressed with the operation and is considering additional MetroCluster deployments, including additional FlexPod that they want to convert to a similar design. They feel very prepared to take on the move of the equipment to the new site when it is completed. In addition, the customer is very happy with the added resiliency of their applications.

Future Work

The initial solution was implemented at the same site since the new facilities had not yet been completed. Since the new site will be relatively close, there is no concern over performance once the equipment is moved.

To automate configuration management, Ansible is being proposed to propagate changes throughout the system. For example, if VLANs are added to one side, it is important to ensure that a change is made to the UCSs, switches, and possibly the upstream firewalls.

Red8 has the technical expertise to assist customers through this process.

Key Technologies Discussed

VMware vSphere Dynamic Resource Scheduler

VMware vSphere® Distributed Resource Scheduler™ (DRS) is a resource scheduling and load balancing solution for vSphere. DRS works on a cluster of ESX hosts and provides resource management capabilities like load balancing and virtual machine (VM) placement. DRS also enforces user-defined resource allocation policies at the cluster level while working with system-level constraints. In vSphere 7, VMware enhanced the DRS logic used in assigning a location to a workload. The new method focuses more on the state of the VM rather than the state of a host—checking every minute to calculate a new score for the efficiency of the VM. As a result, DRS is now better equipped to ensure workloads perform optimally and hosts are used efficiently.

VMware vSphere Metro Storage Cluster (vMSC)

VMware vSphere® Metro Storage Cluster (vMSC) is a specific storage configuration that is commonly referred to as stretched storage clusters or metro storage clusters. In addition, vMSC technology combines replication with array-based clustering. These solutions are typically deployed in environments where the distance between data centers is limited, often metropolitan or campus environments, and in environments where disaster and downtime avoidance are key requirements.

A vMSC infrastructure is often implemented with the goal of reaping the same benefits that vSphere HA clusters provide to a local site in a geographically dispersed model with two data centers in different locations. The architecture is built on the premise of extending what is defined as “local” in terms of network, storage, and compute to enable these subsystems to span geographies, presenting a single and common base infrastructure set of resources to the vSphere cluster at both sites. It, in essence, stretches storage, network, and computing between sites.

The primary benefit of this stretched cluster model is that it enables fully active and workload-balanced data centers to be used to their full potential, and it allows for extremely fast recovery in the event of a host or even full-site failure. The capability of a stretched cluster to provide this active balancing of resources should always be the primary design and implementation goal.

NetApp MetroCluster

NetApp® MetroCluster™ technology combines array-based clustering with synchronous replication to deliver continuous availability, immediately duplicating mission-critical data on a transaction-by-transaction basis. MetroCluster enhances the built-in high availability and non-disruptive operations of NetApp hardware and ONTAP® storage software, providing an additional layer of protection for the entire storage and host environment.

NetApp MetroCluster Tiebreaker software

The NetApp® MetroCluster™ Tiebreaker software checks the reachability of the nodes in a NetApp MetroCluster configuration and the cluster to determine whether a site failure has occurred. The Tiebreaker software triggers an alert under certain conditions. The NetApp MetroCluster Tiebreaker detects direct and indirect failures so that the Tiebreaker does not initiate a switchover if the fabric is intact. The Tiebreaker software resides on a Linux host.

FlexPod

The FlexPod platform, developed by NetApp and Cisco, is a flexible, converged infrastructure solution that delivers pre-validated storage, networking, and server technologies. It is designed to increase IT responsiveness to business demands while reducing the overall computing cost. In addition, FlexPod is an architecture that requires only one support number to get support for either Cisco and/or NetApp without having to perform initial troubleshooting on their own.

VMware vSphere Site Recovery Manager

Site Recovery Manager is an industry-leading disaster recovery (DR) software that delivers automated orchestration of failover and fail-back to minimize downtime. Built-in non-disruptive testing and reporting simplify audits and ensure RTOs are met.

NetApp SnapMirror

NetApp® SnapMirror® replicates data at high speeds over local or wide-area networks to provide high data availability and fast data replication for business-critical applications, including Microsoft Exchange, Microsoft SQL Server, and Oracle, in both virtual and traditional environments.

Replicating data to one or more NetApp storage systems and continually updating the secondary data, data is kept current and remains available whenever you need it. No external replication servers are required to enable this replication technology.

Microsoft Geo-Cluster

VMware vSphere HA

VMware vSphere High Availability delivers the availability required by most applications running in virtual machines, independent of the operating system and applications running in it. High Availability provides uniform, cost-effective failover protection against hardware and operating system outages within your virtualized IT environment.
High Availability allows:

- Monitoring VMware vSphere hosts and virtual machines to detect hardware and guest operating system failures.
- Restarting virtual machines on other vSphere hosts in the cluster without manual intervention when a server outage is detected.
- Reduction of application downtime by automatically restarting virtual machines upon detection of an operating system failure.

Ansible

Ansible is an IT automation tool. It can configure systems, deploy software, and orchestrate more advanced IT tasks such as continuous deployments or zero downtime rolling updates.

It lends itself to strong source code control and management akin to software development processes. As a result, it is often a key component “infrastructure as code” infrastructure management methodology.

Additional Information

For additional information, please contact a Red8 team member at info@red8.com.