

Active-Active Virtualization: A More Cost Effective Approach to Cross Campus Disaster Recovery

Dr Allon Cohen, VP Product Management, SANRAD

Corporations will go a long way to make sure they are able to continually serve their customers. Because the reason and location of a failure can not accurately be predicted, companies will often create a Disaster Recovery (DR) site that will enable them to function if the main site goes down for any reason. A DR site can require a very large investment for the company. To compound this challenge, many companies maintain these sites only to be used for disaster recovery, resulting in very inefficient resource utilization and poor return on investment.

Until recently there was no effective way to use the servers and storage at both sites regularly while still maintaining autonomous recovery capabilities. This has now changed and companies can deploy active-active disaster recovery, fully utilizing their investment in resources at the alternate sites.

Active-Active Disaster Recovery

Figure 1 presents an architecture that allows for active-active disaster recovery. The architecture is an example of a high-availability (synchronous) DR solution working with virtual servers. The figure represents two production sites. Both sites are designed to operate concurrently under normal conditions and to serve as a DR site for each other, should a complete site failure occur at either one. Each site contains a set of physical servers running virtual servers (e.g. Microsoft Hyper-V or vmware HA), and a synchronous campus cluster capable storage (SANRAD's V-STOR).

The key benefit from this architecture is that, although both sites can run concurrently and together serve business users, complete high availability is maintained at both the server level and the storage level between the sites.

Figure 1 – Active-Active Disaster Recovery





Figure 2 illustrates how these two active-active sites serve as backups for each other in the event of a complete failure at one of the sites. When such a disaster occurs, the virtual machines are transferred and reloaded on the servers at the alternate site. At the same time, since the V-STORs have been synchronously mirroring data between these sites, a complete copy of the data is available at the alternate site. Also since the target address for the storage has been transferred to the surviving V-STOR, there is zero interruption to service, and no IT intervention required. As the virtual machines reload at the alternate site, they can immediately reconnect to their data and continue servicing end users automatically.

Figure 2 – Fully Autonomous Failover



Conclusion

Maintaining a DR site is an important element of the business continuity strategy of any corporation. However, up until recently, creating such sites required a large investment in infrastructure and services that remained unused under normal business circumstances. Recent advances in server virtualization and intelligent network storage enable IT departments for the first time to employ DR sites with active-active architectures. With active-active architecture, companies can fully utilize their hardware investment while serving the business continuity requirements of the enterprise.