SIOS
Clusters Your Way.™

Windows Azure IaaS: High Availability and Disaster Recovery with SIOS

By David Bermingham, Microsoft Cluster MVP
Senior Technical Evangelist, SIOS Technology Corp.

Mission Critical SQL Server Needs a Mission Critical Platform

Cloud services such as Windows Azure IaaS open up a realm of possibilities that previously may have been too expensive to implement or too complicated to manage. You can use Windows Azure as an extension of your existing data center, replacing the need to manage an offsite facility. Or, you may choose to eliminate the expense of building your own data center and instead deploy your applications in Windows Azure, taking advantage of the pay-as-you-go model and flexibility to expand capacity on demand. Regardless of how you choose to use Windows Azure, you will have to plan for both disaster recovery and high availability in your deployment.

Although Windows Azure has redundancy as described in Windows Azure Business Continuity Technical Guidance[1], applications such as SQL Server and File Servers still need additional configuration for high availability and disaster recovery. SIOS Technology has data replication and high availability solutions that seamlessly integrate with Windows Azure to facilitate the four different high availability and disaster recovery models described below:

- **Availability within Windows Azure** – Failover Clustering within a Windows Azure deployment.
- **Disaster Recovery for Windows Azure Deployments** – Multisite cluster that extends a Windows Azure Failover Cluster to a cluster node that resides outside of Azure.
- **Windows Azure as a Hot Standby DR Site** – Multisite cluster that extends an on-premise cluster to Windows Azure for automated disaster recovery
- **Windows Azure as a Warm Standby DR Site** – Data replication from on-premise servers into Windows Azure for manual recovery in the event of a disaster.

This document will explain each of the four models mentioned above and provide guidance on how to implement these solutions.

Availability within Windows Azure

Windows Azure IaaS currently has eight geographic regions in which you can deploy your cloud services: West Europe, North Europe, East US, West US, Southeast Asia, East Asia, Japan East, and Japan West. They also have plans to open another datacenter in Brazil in early 2014.[2] Within each datacenter, there is a concept of a fault domain. A fault domain essentially is described as follows:[3]

> *You manage the availability of an application that uses multiple virtual machines by adding the virtual machines to an availability set. Availability sets are directly related to fault domains and update domains. A fault domain in Windows Azure is defined by avoiding single points of failure, like the network switch or power unit of a rack of servers. In fact, a fault domain is closely equivalent to a rack of physical servers. When multiple virtual machines are connected in a cloud service, an availability set places the virtual machines in different fault domains.*

By provisioning your Azure virtual machines (VMs) in the same availability set, you ensure that each VM resides in a different fault domain. With the VMs running in different fault domains, a localized failure will not effect your entire Windows Azure Deployment. For web servers such as IIS, simply putting your servers in different fault domains and enabling load balancing[4] will give you a level of availability. Windows Azure Traffic Manager[5] gives you even more load balancing options, including load balancing across data centers.

When it comes to SQL Server and File Servers, load balancing does not solve the problem of maintaining application availability. For these applications, you must locate redundant VMs in separate Fault Domains and replicate and maintain the application data *identically* between the VMs.

High availability for SQL Server and File Servers traditionally involves Windows Server Failover Clusters (WSFC), which traditionally relies on some sort of shared storage device, such a SAN. But, Windows Azure has no concept of a cluster aware shared storage device, precluding a traditional SAN-based WSFC as an option.

However, you can easily create a SIOS #SANLess cluster in Windows Azure, that supports WSFC and provides high availability for SQL Server, File Servers and any cluster aware application. A sample configuration is shown in Figure 1.

What is a #SANLess Cluster?

You create a #SANLess cluster simply by adding SIOS DataKeeper Cluster Edition[6] software to a Windows Server environment. DataKeeper provides host based, block level, volume replication that integrates with WSFC, allowing you to build failover clusters in Windows Azure IaaS using mirrored disk storage for each VM instead of shared storage.

How do you create #SANLess Cluster?

To create a #SANLess cluster, you follow the same basic steps as you would in creating a traditional cluster, but instead of using a disk witness you use a file share witness. And instead of shared storage, you provision each VM with an additional disk(s), which will be replicated by DataKeeper.

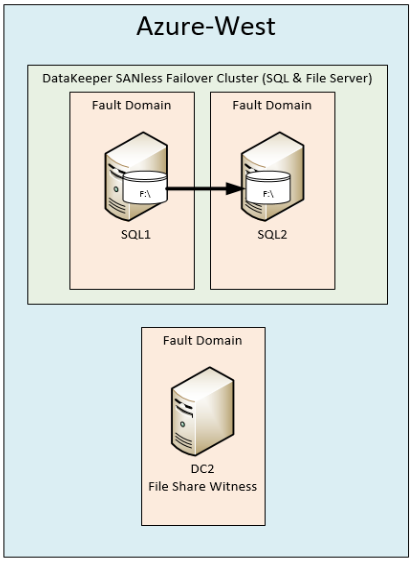

Figure 1 - DataKeeper #SANLess cluster in Windows Azure.

Three Easy Steps to Configuration

To create a mirror, the DataKeeper Cluster Edition software must be installed on each server. The DataKeeper GUI features a simple three-step wizard as shown in Figure 2. The wizard allows you to specify the source and target server, as well as mirror options such as compression, bandwidth throttle, and synchronous or asynchronous mirroring.

Once the mirror is created, a DataKeeper Volume Resource is registered in Available Storage within WSFC as shown in Figure 3 on the next page.

With the DataKeeper Volume in Available Storage, you can create the #SANLess cluster just as if it were a SAN-based cluster. The detailed step-by-step instructions are documented in the article "[Creating a SQL Server 2014 AlwaysOn Failover Cluster (FCI) Instance in Windows Azure IaaS.]"[7] Other cluster resources such as File Servers can be created just as easily by using the DataKeeper Volume Resource in the cluster.

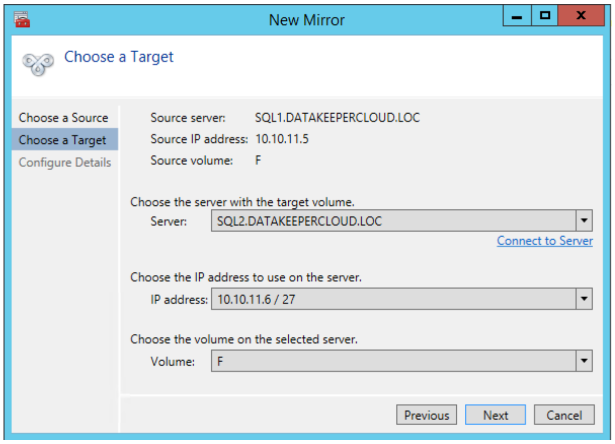

Source server: SQL1.DATAKEEPERCLOUD.LOC
Source IP address: 10.10.11.5
Source volume: F

Choose the server with the target volume.
Server: SQL2.DATAKEEPERCLOUD.LOC

Choose the IP address to use on the server.
IP address: 10.10.11.6 / 27

Choose the volume on the selected server.
Volume: F

Figure 2 - DataKeeper mirror creation wizard.

Multisite Clusters for Providing Disaster Recovery in Windows Azure Deployments

While providing high availability within the Windows Azure Cloud between Fault Domains will protect you from normal hardware failures and other unexpected outages within a Fault Domain, how do you protect from a larger scale outage that might affect an entire Windows Azure geographic region? Natural disaster or human error could potentially render your entire Windows Azure deployment offline. While such an outage is not likely, it would be prudent to plan for an outage just in case.

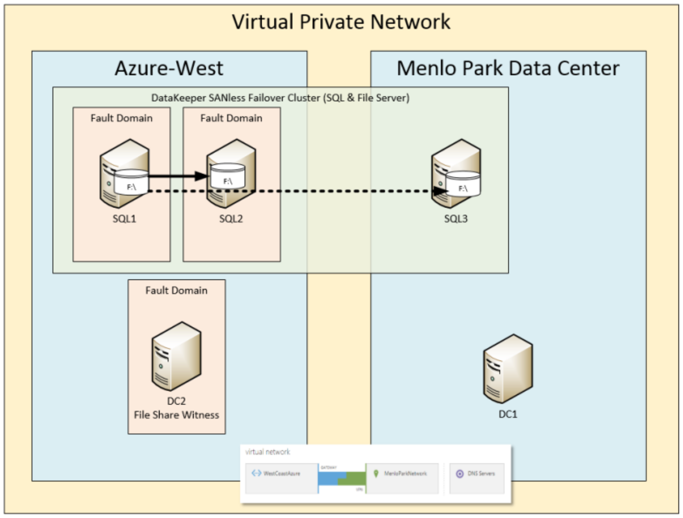

Figure 3 - The DataKeeper Volume Resource in available storage.

One way to plan for such an outage is by building a multisite cluster. As described above, it is possible to build a Failover Cluster in the Windows Azure Cloud. It is also possible to extend that cluster by adding an additional node(s) in an alternate datacenter. A multisite cluster gives you a recovery point objective (RPO) of near zero data loss and a recovery time objective (RTO) of just about one minute.

While Windows Azure currently does not support the creation of a Virtual Network that spans different Azure geographic regions, they do allow you to create a Virtual Network that spans between Azure and your on-premise datacenter using a site-to-site VPN[8] or a dedicated private network using ExpressRoute[9]. Figure 4 shows a three node multisite cluster with two nodes located in Azure West and the third node located in a Menlo Park datacenter.

SIOS DataKeeper Cluster Edition replicates SQL1 and SQL2 synchronously within the Azure network, while at the same time asynchronous replication keeps SQL3 synchronized off site. With new options available in Windows Server 2012 R2 for dynamic quorum, dynamic witness and site preference, it is possible to configure the cluster as shown in Figure 4 and still have automatic failover in the event of a disaster. However, many people will chose to disable failover to SQL3 in this case and instead just use the force quorum method in the event of an actual disaster, especially if their RTO doesn't require failover within minutes.

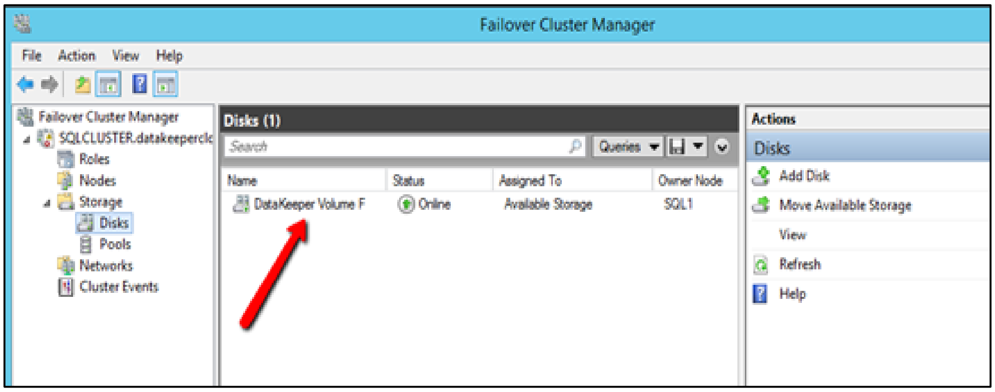

Name	Status	Assigned To	Owner Node
DataKeeper Volume F	Online	Available Storage	SQL1

Figure 4 - Multisite cluster for Azure disaster recovery.

SQL Server Standard Edition Option

In this example, SQL3 is part of the cluster. To build a three-node multisite SQL Server Cluster, you have to use SQL Server 2012/2014 Enterprise Edition, as SQL Server 2012/2014 Standard Edition only supports a two node cluster.

However, if you can tolerate a slightly longer recovery time, you may want to consider saving the licensing cost of Enterprise Edition by using SQL Server Standard Edition with DataKeeper Cluster Edition, which allows you to replicate to a third node outside of the cluster. This gives you added disaster protection, but recovering the database requires mounting the replicated databases and redirecting the clients to the new SQL Server instance.

Windows Azure as a Hot Standby DR Site

If we take the same concept and turn it around such that your on-premise datacenter is your primary datacenter and Windows Azure is your DR site, then you have Windows Azure as a hot standby DR site. This is a very cost effective alternative to building out your own DR site, or renting rack space in a business continuity facility. In this case, the on-premise servers can be traditional SAN based clusters, SANLess clusters or even single servers not currently participating in a cluster.

The objective of having a “hot” standby DR site is to have standby servers up and running in the DR site, with a copy of the most recent application data. In the event of a disaster, recovery should be automatic or “push button”.

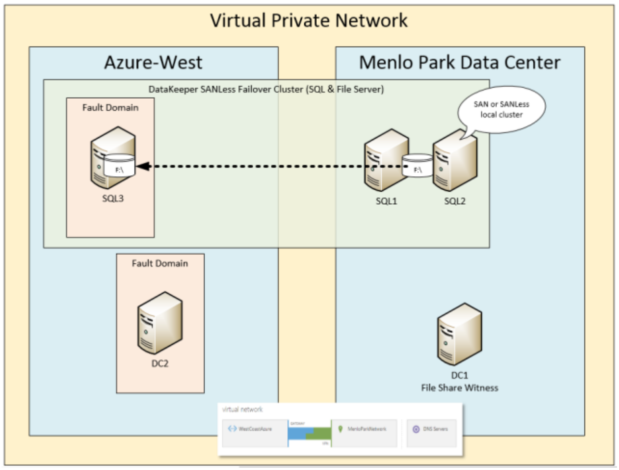

Figure 5 - Azure Windows as a hot DR site.

A multisite cluster as shown in Figure 5 is one of the best ways to implement a hot standby DR site with Windows Azure. DataKeeper Cluster Edition continuously updates the Azure VM(s) using asynchronous replication. In the event of a planned outage, such as a predicted hurricane, applications can be moved to the Azure Cloud using the WSFC interface before potential disaster strikes. In the event of an unexpected disaster, applications can be recovered in Windows Azure manually by forcing the cluster quorum online or in some cases recovery can be automatic, depending upon the quorum configuration. This configuration gives you an excellent RTO and RPO with minimal investment.

Using Windows Azure as a Warm Standby DR Site

If your RTO requirements for disaster are more lenient, you may consider using Azure as a “warm” standby site. See figure 6. In this configuration DataKeeper Standard Edition is used to replicate data from on-premise into Azure. Recovery in the event of a disaster would generally involve starting some VMs that were previously provisioned, but not actively running. The only VMs running in Azure are enough VMs to act as the targets for the DataKeeper replication.

The benefit of a warm standby site over a hot standby site is that most VMs in Azure power off most of the time and will only be running in the event of a disaster, thereby reducing the monthly Windows Azure usage charge. The downside is that recovery will take some man power and time, increasing your RTO. However, the RPO remains the same as in a "hot" standby DR site as DataKeeper keeps the data in sync with asynchronous replication.

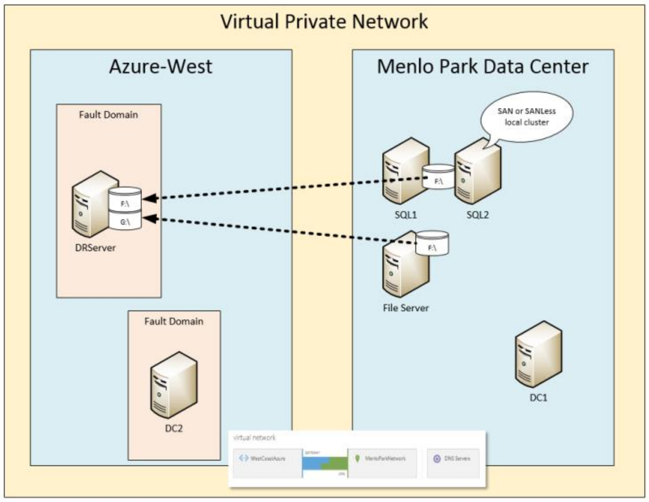

Figure 6 - Azure Windows as a warm DR site.

Summary

Windows Azure IaaS provides a robust platform for hosting mission critical applications. While the infrastructure is in place for high availability and disaster recovery, it is up to the administrator to use the tools at their disposal to build high availability and disaster recovery into any Windows Azure deployment. Windows Azure is also an attractive alternative as a disaster recovery site. With tools like Virtual Networks and site-to-site VPN support, Windows Azure provides you with the infrastructure to make it possible to extend your data center to the cloud. SIOS Technology's DataKeeper Cluster Edition and DataKeeper Standard Edition are the tools that make high availability and disaster recovery in Windows Azure simple to deploy and manage.

References

[1] http://msdn.microsoft.com/en-us/library/windowsazure/hh873027.aspx

[2] http://blogs.msdn.com/b/windowsazure/archive/2013/12/04/expanding-windows-azure-capacity-brazil.aspx

[3] http://www.windowsazure.com/en-us/documentation/articles/manage-availability-virtual-machines/

[4] http://blogs.technet.com/b/cbernier/archive/2013/11/11/load-balance-virtual-machines-in-windows-azure.aspx

[5] http://www.windowsazure.com/en-us/services/traffic-manager/

[6] http://us.sios.com/products/datakeeper-cluster/

[7] http://clusteringformeremortals.com/2014/01/10/creating-a-sql-server-2014-alwayson-failover-cluster-fci-instance-in-windows-azure-iaas-azure-cloud/

[8] http://www.windowsazure.com/en-us/documentation/articles/virtual-networks-create-site-to-site-cross-premises-connectivity/

[9] http://www.windowsazure.com/en-us/services/expressroute/