

SharePoint Disaster Recovery to Microsoft Azure

Example architectures for building a recovery environment in Azure

Overview

The disaster recovery environment for an on-premises SharePoint 2013 farm can be hosted in Azure.

- Azure Infrastructure Services provides a secondary datacenter.
- Pay only for the resources you use.
- Small recovery farms can be scaled out after a disaster to meet scale and capacity targets.

The recovery farm in Azure is configured as identically as possible to the production on-premises farm.

- Same representation of server roles.
- Same configuration of customizations.
- Same configuration of search components (these can be on a smaller version of the production farm).

Log shipping and Distributed File System Replication (DFSR) are used to copy database backups and transaction logs to the Microsoft Azure farm.

- DFSR is used to transfer logs from the production environment to the recovery environment. In a WAN scenario DFSR is more efficient than shipping the logs directly to the secondary server in Microsoft Azure.
- Logs are replayed to the Microsoft Azure-based SQL Server computers.
- Log-shipped databases are not attached to the farm until a recovery exercise is performed.

Failover procedures:

- Stop log shipping.
- Stop accepting traffic to the primary farm.
- Replay the final transaction logs.
- Attach the content databases to the farm.
- Start a full crawl.
- Restore service applications from the replicated services databases.

Recovery objectives provided by this solution include:

- ✓ Sites and content
- ✓ Search (re-crawled, no search history)
- ✓ Services

Additional items that can be addressed by Microsoft Consulting Services or a partner:

- Synchronizing custom farm solutions
- Connections to data sources on-premises (BDC and search content sources)
- Search restore scenarios
- Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO)

Cold standby environments take longer to start but are less expensive

- The farm is fully built, but the virtual machines are stopped after the farm is created. You only pay processing costs when the virtual machines are running, but storage and network data transfer costs apply.
- In the event of a disaster, all the farm virtual machines are started and patched.
- Backups and transaction logs are applied to the farm databases.

Additional procedures for cold standby environments

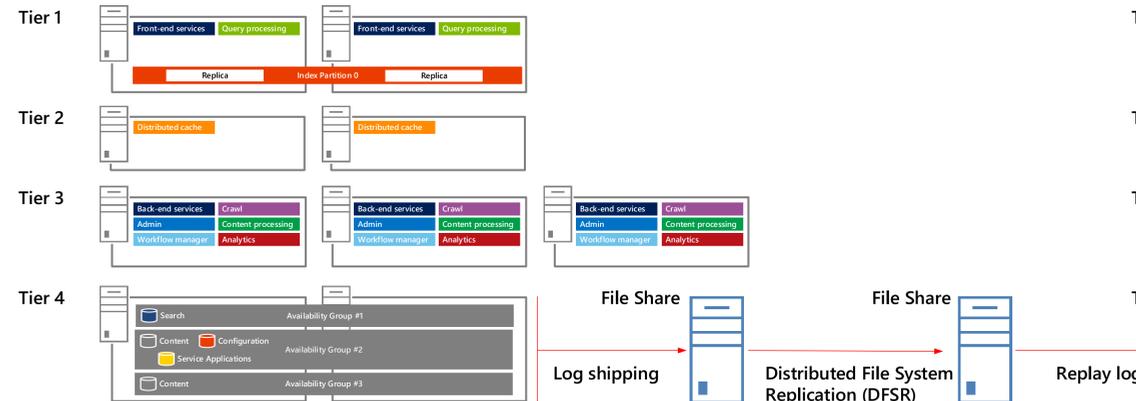
- Turn on virtual machines regularly to patch, update, and verify the environment.
- Run procedures to refresh DNS and IP addresses.
- Setup SQL Server AlwaysOn after a failover.



On-premises environment

Production environment

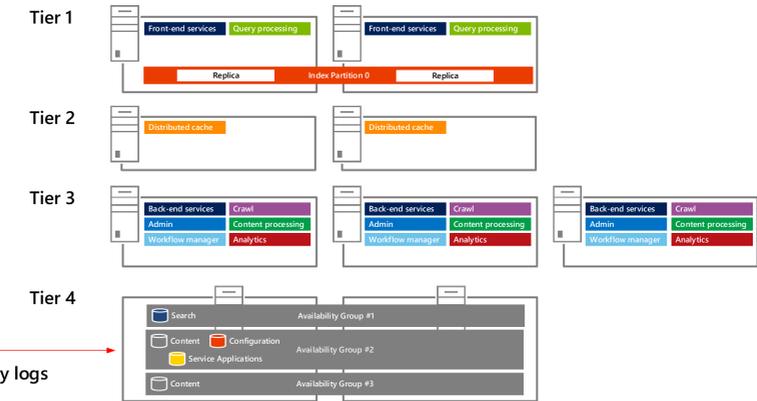
Live production environment



Microsoft Azure recovery environment

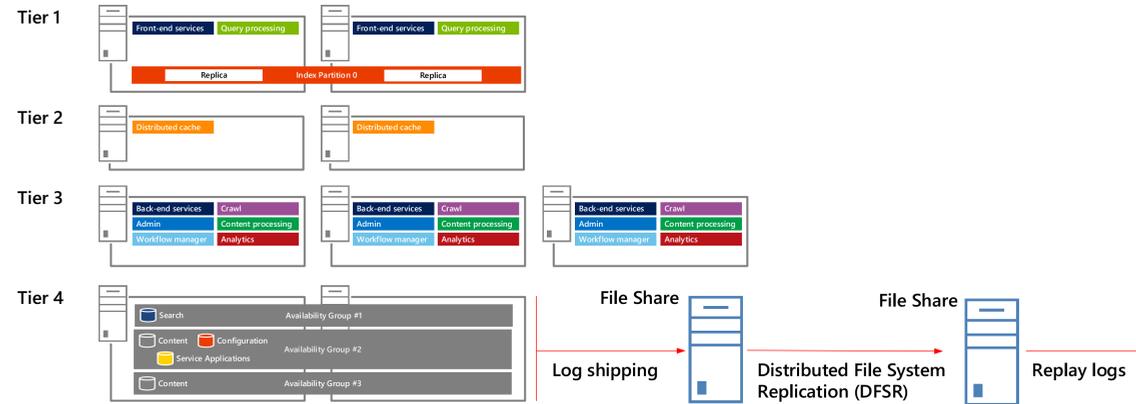
Warm standby environment

Running VMs



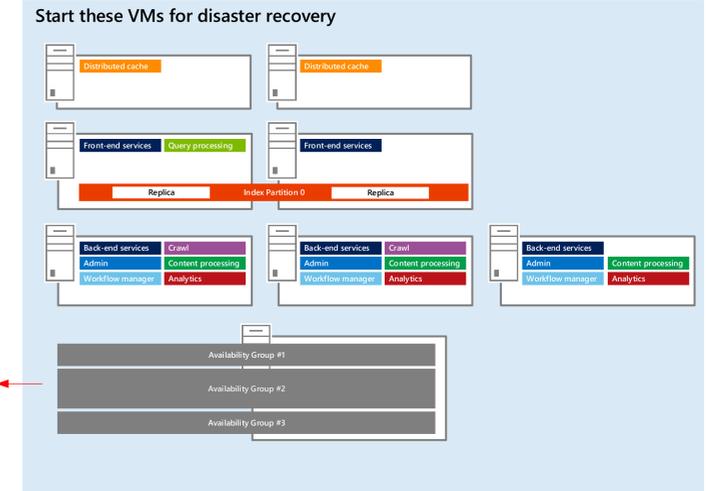
Production environment

Live production environment



Cold standby

Running VMs

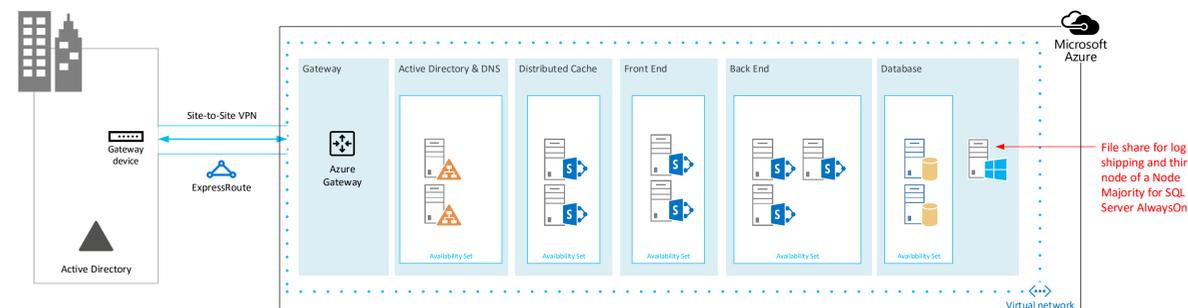


SharePoint recovery environment in Microsoft Azure

Design and build the failover environment in Azure

- Create a cross-premises virtual network in Azure.
- Connect the on-premises network with the virtual network in Azure with a site-to-site VPN or ExpressRoute connection. This connection uses a dynamic gateway in Azure.
- Deploy one or more Windows Server AD domain controllers (DCs) to the Azure virtual network and configure these to work with your on-premises domain. These DCs are catalog servers.
- Deploy the tiers of the SharePoint farm on different subnets and in different Azure availability sets.
- Deploy the SharePoint farm plus a file server to host file shares.
- Setup log shipping and DFSR between the on-premises environment and the Azure-based recovery environment.

For detailed instructions to deploy a warm standby environment, see [SharePoint Server 2013 Disaster Recovery in Microsoft Azure](#).



Build the Windows Server Active Directory hybrid environment

The configuration of Windows Server Active Directory (AD) for this solution constitutes a hybrid deployment scenario in which Windows Server AD is partly deployed on-premises and partly deployed on Azure virtual machines.

This reference architecture includes two virtual machines configured as domain controllers. Each is configured as follows:

- **Size** — Small.
- **Operating system** — Windows Server 2012 R2.
- **Role** — Windows Server AD domain controller designated as a global catalog server. This configuration reduces egress traffic across the cross-premises connection. In a multi-domain environment with high rates of change, configure domain controllers on-premises to not synchronize with the global catalog servers in Azure.
- **Data disks** — Place the Windows Server AD database, logs, and SYSVOL on Microsoft Azure data disks. Do not place these on the operating system disk or the temporary disk provided by Azure. This is important.
- **Role** — Install and configure Windows DNS on the domain controllers.
- **IP addresses** — Use dynamic IP addresses.