

# Datacenter and Service Recovery

How Microsoft services recover from a datacenter loss



Published August 2017

## Disclaimer

*This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.*

*This document is provided “as-is.” Information and views expressed in this document, including URL and other Internet website references, may change without notice. Customers reading this document bear the risk of using it.*

*This document does not provide customers with any legal rights to any intellectual property in any Microsoft product. Customers may copy and use this document for their internal, reference purposes.*

*The information contained in this document must not be construed as legal advice. Customers must seek their own legal counsel for advice on compliance with regulatory requirements impacting their organization.*

*Some examples depicted herein are provided for illustration only and are fictitious. No real association or connection is intended or should be inferred.*

*NOTE: Certain recommendations in this document may result in increased data, network, or compute resource usage, and may increase a customer’s license or subscription costs.*

*Version 1.1*

*© 2017 Microsoft. All rights reserved.*

## Acknowledgements

### Author

Robert Arco

### Contributors and Reviewers

Eric Tierling

Katie Jackson

Alessandra Reyes

Jeffrey Gallucci

Steve Wacker (Wadeware LLC)

Table of contents

- Introduction ..... 4**
  - Overview ..... 4**
  - Shared responsibilities between Microsoft and customers ..... 4**
- 1. Datacenter recovery ..... 5**
  - Traditional datacenter failover ..... 5**
    - Primary datacenter recovery to a cold site location ..... 5
    - Primary datacenter recovery to a warm site or hot site ..... 5
    - Active-Active, parallel processing, and parallel sysplex ..... 5
    - Sample configurations for traditional datacenter recovery ..... 6
  - Microsoft service and regional level recovery ..... 6**
  - Azure paired regions ..... 6**
  - Regional recovery vs Service recovery ..... 7**
- 2. Microsoft Azure BCDR programs ..... 9**
  - Azure services certifications ..... 9**
- 3. Key information on Azure service recovery ..... 10**
  - Essential elements of service recovery architected by Microsoft services ..... 10**
  - Guidance for customers ..... 10**
- Conclusion ..... 11**

## Introduction

One of the most common questions Microsoft receives from customers and auditors is “Tell us how your datacenters failover in the event of catastrophic loss”. The question itself is dated because Microsoft does not use the concept of datacenter failover; instead, failover is relegated first to the individual services and catastrophic loss is relegated to paired regions for recovery.

In the traditional IT operations model, when organizations began addressing continuity of service for their technology and data, the strategy was to recover a lost primary datacenter to an alternate datacenter location in total. Various iterations of that concept were implemented that addressed complete technology and data restoration. In some cases, these iterations were automated (at high cost) and in others, restoration entailed recovery that could span days because of slow data recovery methods such as tape restore. In addition, this approach called for two complete “like for like” datacenters, required all services to recover to the alternate datacenter, and limited recovery to complete loss scenarios.

To support continuity of services, Microsoft chose to architect recovery at both the service level and region level, rather than the datacenter level, in order to support a broader spectrum of recovery. Loss of a datacenter is not the only event that can affect services, and by focusing on the service level Microsoft can help ensure recoverability of services for smaller, limited events. In other words, many individual services have the ability to fail to another location without a full datacenter recovery, and those services may be run from a variety of locations.

## Overview

This paper explains how Microsoft achieves service reliability and recoverability regardless of impact and scope of outage. It should serve to answer the question in modern terms of the service failover versus datacenter failover concepts. Specific failover information for individual services can be found under separate references.

## Shared responsibilities between Microsoft and customers

Understanding how cloud service providers (CSPs) such as Microsoft share responsibility with customers to meet security, privacy, and compliance requirements is an essential part of cloud-based computing.

When adopting Microsoft cloud services, it is important to remember that some security, privacy, and compliance needs are the responsibility of the customer, some are the responsibility of Microsoft, and some are shared. Read the [Shared responsibility in cloud computing](#) white paper to learn more about the responsibility for each party in a cloud-based solution.

# 1. Datacenter recovery

In the traditional IT operations model, datacenter recovery leveraged various strategies to recover from catastrophic loss. In most cases, these strategies were limited to support of major event recovery and did not necessarily address service resiliency in terms of availability and reliability at the single service level. In-place repair for services was the norm for individualized events. This restriction relegated service recovery to only total-loss scenarios.

## Traditional datacenter failover

Traditional datacenter failover scenarios work as follows.

### Primary datacenter recovery to a cold site location

Because the cold site is not active on a regular basis, it requires activation, configuration, recovery (network, service, and data), and often travel for supporting workforce. The cold site could be owned by the primary client or be contracted to a third-party service. This strategy is considered the slowest recovery option and, short of a major datacenter impact, individual service impacts could only be addressed by in-place repair.

### Primary datacenter recovery to a warm site or hot site

A warm or hot site scenario reduces recovery time by having prepared configurations but still requires steps to activate networks, recover service availability, and in many cases data recovery that may have experienced significant recovery point impacts. The sites could be self-owned or third-party supported and, like the cold site, only used to support complete datacenter failure and did not provide service level recovery for limited events.

### Active-Active, parallel processing, and parallel sysplex

These configurations are closest to the current Microsoft scenario. However, besides being costly, they are generally limited to an A-B structure. The structures need to be far enough away from each other (geographically dispersed) to remediate regional impacts but close enough that transactional services are not impacted by physics (speed of light) causing backlog or dropped transactions and data.

## Sample configurations for traditional datacenter recovery

In the traditional datacenter recovery model, sample configurations usually look like this:

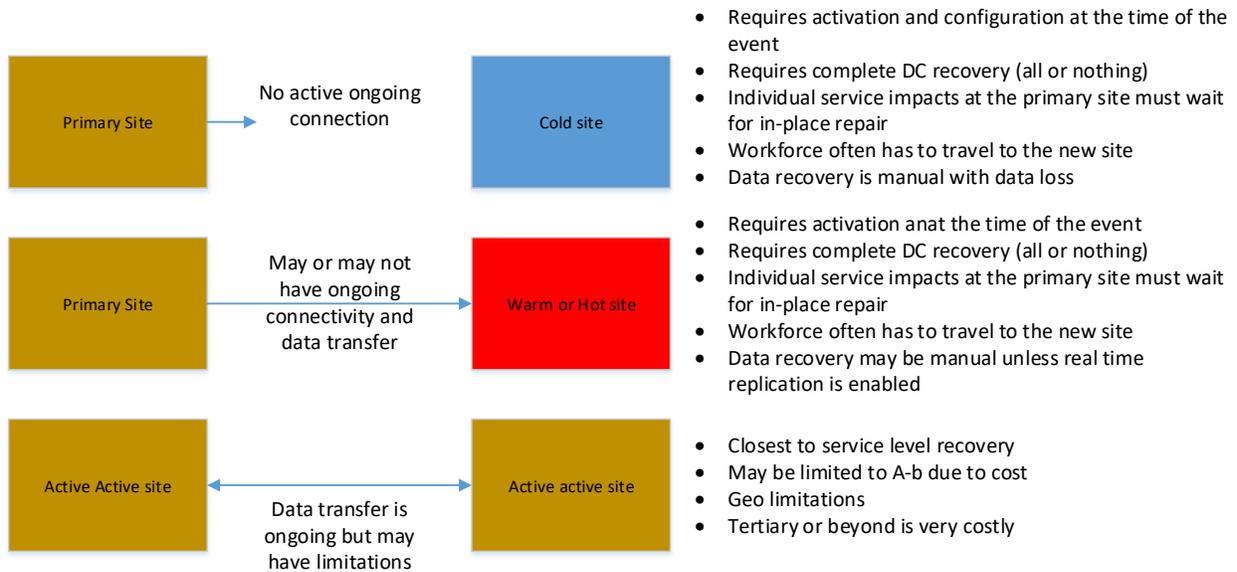


Figure 1. Sample configurations for traditional datacenter recovery

## Microsoft service and regional level recovery

For its cloud services, Microsoft designs and supports recovery at both service level and region level. However, the approximately 100 datacenters around the globe are not simply configured into A & B pairs; instead, each individual service designs, selects, and tests recoverability to a variety of sites. This approach allows many Microsoft services to recover at alternate locations that do not require a full datacenter recovery. Services can remain highly available if the recovery of down services is decoupled from full datacenter loss. This strategy eliminates the need for services to wait for in-place repair.

Although service level recovery is possible for many of its systems, Microsoft also addresses total loss scenarios via regional recovery that uses paired regions. Unlike a traditional datacenter, a region may be made up of several datacenters providing high availability. If a region is truly lost due to a regional event, the recovery is performed to a paired region. In most cases, all necessary cloud services are already in place (either by Microsoft or customer configuration using Azure building blocks) and critical data has been geo-replicated real time, eliminating the need for slower restoration methods.

## Azure paired regions

Azure operates in multiple geographies around the world. An Azure *geography* is a defined area of the world that contains at least one Azure region. An Azure *region* is an area within a geography containing one or more datacenters.

Each Azure region is paired with another region within the same geography, together making a regional pair (the exception is Brazil South, which is paired with a region outside its geography). More information on Azure paired regions can be found at [Azure.com](https://azure.com).

Geography	Paired regions	
North America	North Central US	South Central US
North America	East US	West US
North America	US East 2	US Central
North America	West US 2	West Central US
Europe	North Europe	West Europe
Asia	South East Asia	East Asia
China	East China	North China
Japan	Japan East	Japan West
Brazil	Brazil South <sup>1</sup>	South Central US
Australia	Australia East	Australia Southeast
US Government	US Government Iowa	US Government Virginia
India	Central India	South India
Canada	Canada Central	Canada East
UK	UK West	UK South

Table 1. Mapping of Azure regional pairs

## Regional recovery vs Service recovery

Although datacenter or regional catastrophic impacts are extreme, they are of a low probability.

<sup>1</sup> Brazil South is unique because it is paired with a region outside of its own geography. Brazil South's secondary region is South Central US, but South Central US's secondary region is not Brazil South.

Individual service impacts caused by hardware, network, power, and so on have a higher probability of occurring. By designing recovery and resilience at the core service level, Microsoft customers can be assured of better service availability and reliability. In addition, Microsoft services address service resiliency well beyond the concept of the traditional IT operations model for disaster recovery.

As previously described, Regional recovery is limited to only the most extreme loss scenarios that necessitate a full exit of the primary region location. The difference in Service recovery is that individual services can failover to alternate regions for limited losses specific to that service. They can then support the primary region from the alternate region location.

Although this may seem like the typical Active-Active datacenter concept, it differs by allowing many of the services to choose from a variety of recovery locations, not all of which are limited to the paired region. Services design the best fit for recovery locations and often choose several locations for critical service availability. An example would be five services residing in a U.S. West region. Three of those services may be set up to failover to U.S. Central while two others are designed to fail to U.S. East. What this shows is that at the service level, they are not locked in to a singular recovery location.

Another factor of Microsoft service level recovery is the ability for services to leverage multiple instances of a service that serve specific "closest" areas but fully support each other for down locations. A service may deploy in 8 locations globally to provide the best availability; however, it would be configured to provide next-closest-location support for Microsoft customers if a single location goes offline.

## 2. Microsoft Azure BCDR programs

Microsoft takes business continuity and service resiliency seriously on behalf of its customers. The BCDR (business continuity and disaster recovery) and service resiliency programs are an inherent part of the Microsoft culture of providing reliable services.

Azure BCDR programs have the highest number of industry and government certifications of any cloud service provider. In September 2016 Microsoft Azure achieved [ISO 22301 certification](#); Microsoft is the first hyper-scale cloud provider to achieve this important certification for business continuity management. It helps ensuring that Azure applications are backed by the highest standard for business continuity and disaster preparedness.

### Azure services certifications

Helping organizations to comply with national, regional, and industry-specific requirements that govern the collection and use of individuals' data, Microsoft offers the most comprehensive set of certifications and attestations of any cloud service provider. At the [Microsoft Trust Center](#) you can find out more about Microsoft compliance efforts, while the Microsoft [Service Trust Portal](#) features a wealth of additional and more detailed compliance information.

<b>Global</b>	<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> ISO 27001:2013</li> <li><input checked="" type="checkbox"/> ISO 27017:2015</li> <li><input checked="" type="checkbox"/> ISO 27018:2014</li> </ul>	<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> ISO 22301:2012</li> <li><input checked="" type="checkbox"/> ISO 9001:2015</li> </ul>	<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> SOC 1 Type 2</li> <li><input checked="" type="checkbox"/> SOC 2 Type 2</li> <li><input checked="" type="checkbox"/> SOC 3</li> </ul>	<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> CSA STAR Certification</li> <li><input checked="" type="checkbox"/> CSA STAR Attestation</li> <li><input checked="" type="checkbox"/> CSA STAR Self-Assessment</li> </ul>
<b>US Gov</b>	<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> FedRAMP High</li> <li><input checked="" type="checkbox"/> FedRAMP Moderate</li> </ul>	<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> DoD DISA SRG Level 5</li> <li><input checked="" type="checkbox"/> DoD DISA SRG Level 4</li> <li><input checked="" type="checkbox"/> DoD DISA SRG Level 2</li> <li><input checked="" type="checkbox"/> DFARS</li> </ul>	<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> DoE 10 CFR Part 810</li> <li><input checked="" type="checkbox"/> NIST SP 800-171</li> <li><input checked="" type="checkbox"/> FIPS 140-2</li> <li><input checked="" type="checkbox"/> Section 508 VPATs</li> </ul>	<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> NIST CSF</li> <li><input checked="" type="checkbox"/> ITAR</li> <li><input checked="" type="checkbox"/> CJIS</li> <li><input checked="" type="checkbox"/> IRS 1075</li> </ul>
<b>Industry</b>	<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> PCI DSS Level 1</li> <li><input checked="" type="checkbox"/> GLBA</li> <li><input checked="" type="checkbox"/> FFIEC</li> <li><input checked="" type="checkbox"/> Shared Assessments</li> <li><input checked="" type="checkbox"/> FISC (Japan)</li> </ul>	<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> HIPAA BAA</li> <li><input checked="" type="checkbox"/> HITRUST</li> <li><input checked="" type="checkbox"/> 21 CFR Part 11 (GxP)</li> <li><input checked="" type="checkbox"/> MARS-E</li> </ul>	<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> IG Toolkit (UK)</li> <li><input checked="" type="checkbox"/> NEN 7510:2011 (Netherlands)</li> <li><input checked="" type="checkbox"/> FERPA</li> </ul>	<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> CDSA</li> <li><input checked="" type="checkbox"/> MPAA</li> <li><input checked="" type="checkbox"/> FACT (UK)</li> </ul>
<b>Regional</b>	<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> Argentina PDPA</li> <li><input checked="" type="checkbox"/> Australia CCSL / IRAP</li> <li><input checked="" type="checkbox"/> Canada Privacy Laws</li> <li><input checked="" type="checkbox"/> China GB 18030:2005</li> <li><input checked="" type="checkbox"/> China DJCP (MLPS) Level 3</li> </ul>	<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> China TRUCS / CCCPPF</li> <li><input checked="" type="checkbox"/> EU ENISA IAF</li> <li><input checked="" type="checkbox"/> EU Model Clauses</li> <li><input checked="" type="checkbox"/> EU – US Privacy Shield</li> <li><input checked="" type="checkbox"/> Germany IT-Grundschutz workbook</li> </ul>	<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> India MeitY</li> <li><input checked="" type="checkbox"/> Japan CS Mark Gold</li> <li><input checked="" type="checkbox"/> Japan My Number Act</li> <li><input checked="" type="checkbox"/> Netherlands BIR 2012</li> <li><input checked="" type="checkbox"/> New Zealand Gov CIO Fwk</li> </ul>	<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> Singapore MTCS Level 3</li> <li><input checked="" type="checkbox"/> Spain ENS</li> <li><input checked="" type="checkbox"/> Spain DPA</li> <li><input checked="" type="checkbox"/> UK G-Cloud</li> <li><input checked="" type="checkbox"/> UK Cyber Essentials Plus</li> </ul>

Figure 2. Microsoft Azure has more compliance certifications and attestations than any cloud service provider (as of June 2017)

### 3. Key information on Azure service recovery

Microsoft and Azure cloud services use a variety of strategies to help ensuring service resiliency at the service level. As this paper emphasizes, traditional datacenter recovery is a thing of the past for Microsoft Azure cloud services because of its limitations and because of Microsoft's focused strategy on service availability and reliability for its customers.

#### Essential elements of service recovery architected by Microsoft services

Azure datacenters are built in regions with multiple datacenters per region. Azure is designed to provide high availability within a region in the following ways:

- Azure is architected to provide resiliency for failures (hard drives, servers, network, fault domains) and has self-healing capabilities.
- Regions are physically and logically partitioned to contain failures.
- Regions are in separate fault zones (not susceptible to the same regional hazards) and are physically and logically isolated.
- Cross-region recovery for services is tested frequently for multi-region recovery as part of the Microsoft BCDR program.
- Regional-only services provide customer enablement documentation steps for cross-region availability (customer responsible).
- Datacenters are built and equipped "like for like" so recovery is not constrained when it selects its alternate operational space.<sup>2</sup>
- Recovery locations meet all the needs of the primary location, including security, privacy, capability, and throughput.

#### Guidance for customers

Depending on the service, customers can do the following:

- For regional only services – Run resources across regions to avoid datacenter failures. This configuration is recommended for resources like VMs, where running an active/active or active/passive deployment helps mitigate failures.
- Use geo-redundant resources that automatically replicate data/state across 2 regions. This configuration is recommended for geo-redundant storage (GRS), where the customer's data is automatically replicated across regions within the same geographical area. If a datacenter fails, the customer can continue accessing their data in the secondary region.

---

<sup>2</sup> Not every datacenter is the same but predetermined failover sites meet all the needs of the services.

- Also, customers should review [Disaster recovery and high availability for applications built on Microsoft Azure](#) for more details on how Microsoft thinks about high availability and disaster recovery. The article also has links to some good information that might be useful for this discussion.

## Conclusion

By leveraging service level recovery coupled with region pairing, Microsoft services maintain a high level of reliability and recoverability. By freeing the bounds of full datacenter failover requirements, Azure services and Microsoft services in general have a better ability to handle events at varying scales and not just the datacenter loss scenario.

For more information and details on Azure Services, please visit <https://azure.microsoft.com>.