

Be more resilient and drive innovation

Minimize disruption to mission critical systems and proceed with confidence



Design and operate mission-critical systems with confidence.

	<p>Single VM Locally Redundant Storage (LRS)* Improve the availability of single-instance VMs by using premium/ultra disks to qualify for an availability SLA.</p>	99.9% SLA
	<p>Local redundancies Locally Redundant Storage (LRS) with Azure Managed Disks* Protect against failures with redundancy within a single datacenter in the event of hardware malfunctions or software update cycles.</p>	99.9% SLA
	<p>Zonal redundancies Zone-Redundant Storage (ZRS) Protect against datacenter failures through redundancy within a single region in the event of power, cooling, or networking issues.</p>	99.9% SLA
	<p>Regional redundancies Geo-Redundant Storage (GRS)* Protect against entire-region failures with redundancy beyond a single region in the event of a tornado, earthquake, or other large-scale disaster.</p>	Industry Leading

Custom design your applications to optimize cost.

Many applications do not need 100% high availability; being aware of this at the beginning of planning can help to optimize costs during non-critical periods. Design your recovery strategy to protect against zonal, regional, and application-level failure. The following checklist covers the scope of resiliency planning.

	<p>Define your resiliency requirements based on business needs.</p>	
	<p>Design for resiliency. Start with an architecture that follows proven practices and then identify the possible failure points.</p>	
	<p>Implement strategies to detect and recover from failures.</p>	
	<p>Test the implementation by simulating faults and triggering forced failovers.</p>	
	<p>Deploy the application into production using a reliable and repeatable process.</p>	
	<p>Monitor the application to detect failures. By monitoring the system, you can gauge the health of the application and respond to incidents if necessary.</p>	
	<p>Respond if there are failures that require manual interventions.</p>	

Take advantage of built-in features.

<p>High availability Maintain acceptable continuous performance despite temporary failure in services, hardware, or datacenters—as well as fluctuation in load—using Azure Availability Zones and availability sets.</p>	<p>Disaster recovery Protect against the loss of an entire region through asynchronous replication for failover of virtual machines and data using services like geo-redundant storage and Azure Site Recovery.</p>	<p>Backup and restore Replicate virtual machines and data to one or more regions using Azure Backup, and conduct self-service recoveries of Azure VMs or disks from a secondary region during an outage.</p>

Cloud services are a shared responsibility.

In the traditional on-premises model, the entire responsibility of managing falls on you. With IaaS, the cloud service provider is responsible for the core infrastructure resiliency, including storage, networking, and compute. As you move from IaaS to PaaS and then to SaaS, you'll find that you're responsible for less and the cloud service provider is responsible for more.

■ Managed by customer
■ *Managed by vendor*

On-premises	Infrastructure (IaaS)	Platform (PaaS)	Software (SaaS)
Applications	Applications	Applications	<i>Applications</i>
Data	Data	Data	<i>Data</i>
Runtime	Runtime	<i>Runtime</i>	<i>Runtime</i>
Middleware	Middleware	<i>Middleware</i>	<i>Middleware</i>
O/S	<i>O/S</i>	<i>O/S</i>	<i>O/S</i>
Visualization	<i>Visualization</i>	<i>Visualization</i>	<i>Visualization</i>
Servers	<i>Servers</i>	<i>Servers</i>	<i>Servers</i>
Storage	<i>Storage</i>	<i>Storage</i>	<i>Storage</i>
Networking	<i>Networking</i>	<i>Networking</i>	<i>Networking</i>

*Optional: Azure Backup