



New Zealand cloud region partner playbook

December 2023



The upcoming launch of the New Zealand cloud region represents a tremendous opportunity for partners and customers. By supporting customers to store and process their data on shore, you can help them to meet their data residency and compliance obligations, to reduce the latency of their workloads, and to meet their sustainability commitments.

This document provides some specific guidance for New Zealand North, which is the name of the New Zealand cloud region.

The guidance is centred around two key parts:

1. Three common scenarios, which collectively represent many of the customer use cases for New Zealand North.
2. A set of guidance about technical topics to help customers to plan their use of the new region.

Intended audience

This is a technical document, and it's intended for solution architects and other technical stakeholders who work with customers to plan their use of the New Zealand North region.

Table of contents

What is New Zealand North?	4
Scenario 1: Green field, new deployment	6
Scenario 2: Brown field, multi-region	7
Scenario 3: Brown field, repatriation	8
Service availability and timelines	11
Availability zones, zone redundancy, and zonal services	13
Connectivity to on-premises environments	15
Landing zones	17
Security considerations for New Zealand North	18
Data residency	19
Compliance and regulatory standards	20



What is New Zealand North?

Many Microsoft cloud services will be available from the New Zealand cloud region after its launch. Microsoft Azure, as well as most of the Microsoft 365 and Microsoft Dynamics services, are planned to be available.

By using the New Zealand cloud region, customers gain several benefits:

- **Data residency:** Customers can honour data residency commitments by storing and processing data within New Zealand's borders. By using the New Zealand cloud region, customers can meet their regulatory and cultural data requirements.
- **Security and compliance:** Customers can protect their business with local, regional, and global security offerings. Additionally, organizations who work with sensitive data often have heightened security and compliance requirements.
- **Latency:** Customers can increase their efficiency with faster network connectivity and lower latency between local datacentre regions and the cloud. Customers can build a hybrid environment that includes the local region, which reduces connection latency due to distance or quality of the networks that exist between datacentres.
- **Sustainability:** By using the New Zealand cloud region, customers can save energy costs

and reduce the environmental impact of their operations. Customers can take advantage of Microsoft's next-generation innovations and local partnerships with energy providers and other organisations, and they can reduce their energy usage and carbon emission.

It's important to know that the New Zealand North region is more than just a datacentre. The region includes three availability zones, which are distinct physical locations that are close enough to have low-latency connections to other availability zones, but are far enough apart to reduce the likelihood that more than one will be affected by local outages or weather. Availability zones have independent power, cooling, and networking infrastructure. They're designed so that if one zone experiences an outage, then regional services, capacity, and high availability are supported by the remaining zones.

In addition, every region includes extensive networking infrastructure that enables both public and private connectivity from customers' environments to the Microsoft cloud. The New Zealand North region is connected to Microsoft's global wide area network (WAN), which provides high-bandwidth, low-latency connectivity to other Azure regions internationally.

Common customer scenarios

Scenario 1:

Green field, new deployment

Customers who are new to Azure are likely to deploy workloads directly into New Zealand North.

Key questions to ask


When a customer is deploying a new Azure estate to New Zealand North, here are some questions that are useful to ask:

- **What are their timelines?** Services in New Zealand North will become available in stages. If a customer has a hard deadline, they need to ensure that it fits within the launch timelines and allows for contingency. For more information, see [Service availability and timelines](#).
- **Which services will they use?** New Zealand North will support all foundational and mainstream Azure services, with strategic services deployed on an ad hoc basis based on business cases. Consider whether a customer needs any specific services that might not be available in the region. For more information, see [Service availability and timelines](#). Similarly, consider whether the customer uses third-party services from the Azure Marketplace, which also might not be available in new regions until the ISV has enabled the service for the region.
- **How is the customer likely to grow their use of services within the region?** We need to track the estimated growth in service usage by customers to support region capacity planning.
- **What are the customer's requirements for high availability and disaster recovery?** Confirm they understand that New Zealand North has three independent availability zones, and that they've [fully considered how to use availability zones](#) in their solution architecture.

- **Is the customer planning to use multiple regions?** For example, a customer might plan a global expansion of their business, or to support international employees or business partners. Such a strategy might result in the use of other Azure regions in addition to New Zealand North. Similarly, if customers plan to use strategic services, they might need to run across multiple Azure regions because New Zealand North won't have all services available.

Key actions

- **Understand the services and SKUs** that the customer intends to use, and verify they are supported in the region. For more information, see [Service availability and timelines](#).
- **Plan the HA and DR strategy**, and confirm the customer is comfortable architecting with availability zones. For more information, see [Availability zones, zone redundancy, and zonal services](#).
- **Decide on a connectivity approach**. For more information, see [Connectivity to on-premises environments](#).
- **Establish a landing zone and use New Zealand North for regional components**. For more information, see [Landing zones](#).



Scenario 2:

Brown field, multi-region

Many customers have an existing Azure estate that uses other regions. When they plan a new component or workload, they can deploy to New Zealand North.

Key questions to ask


When a customer already has an Azure estate in another region and talks about deploying a new Azure workload to New Zealand North, here are some questions that are useful to ask:

- **What are their timelines?** Services in New Zealand North will become available in stages. If a customer has a hard deadline, they need to ensure that it fits within the launch timelines and allows for contingency. For more information, see [Service availability and timelines](#).
- **Which services will they use?** New Zealand North will support all foundational and mainstream Azure services, with strategic services deployed on an ad hoc basis based on business cases. Consider whether a customer needs any specific services that might not be available in the region. For more information, see [Service availability and timelines](#). Similarly, consider whether the customer uses third-party services from the Azure Marketplace, which also might not be available in new regions until the ISV has enabled the service for the region.
- **What are the customer's requirements for high availability and disaster recovery?** Confirm they understand how New Zealand North's availability

zone-based architecture works, and that they've [fully considered how to use availability zones](#) in their solution architecture.

Key actions

- **Understand the services and SKUs** that the customer intends to use, and verify they are supported in the region. For more information, see [Service availability and timelines](#).
- **Plan the HA and DR strategy**, and confirm the customer is comfortable architecting with availability zones. For more information, see [Availability zones, zone redundancy, and zonal services](#).
- **Plan how connectivity to the new region will work**. For example, the customer might deploy additional ExpressRoute circuits, upgrade an existing circuit, or deploy a new site-to-site VPN. For more information, see [Connectivity to on-premises environments](#). The customer might also need to plan cross-region connectivity.
- **Plan how the customer's landing zone will be extended to the new region**. For more information, see [Landing zones](#).



Scenario 3:

Brown field, repatriation

Many customers who already have workloads in another region are interested in repatriating those workloads to New Zealand North when it's available. Generally, the workloads are deployed to Australia East or Australia Southeast, but they might be deployed to other regions and the same considerations apply.

Commercial drivers

Repatriating or moving a workload typically results in no change in consumption, so Microsoft is unlikely to invest in or get heavily involved in these activities. There are some exceptions, such as helping larger customers to unlock a new workload by moving a related existing workload.

However, from a broader perspective, repatriation is an important activity that we encourage for a few reasons:

- We know from prior experience that customers who run workloads in their own region are more likely to deploy future workloads.
- By moving workloads onshore, it helps to justify ongoing investment in the New Zealand North region by demonstrating usage and growth.
- For some solutions, customers are likely to gain performance benefits by reducing network latency.
- If other regions have capacity constraints in future, that country's customers are likely to be prioritised for capacity allocation over customers in other countries.

Key questions to ask

When a customer talks about repatriating some or all of their Azure estate to New Zealand North, here are some questions that are useful to ask:

- **What's their motivation?** They might be interested in service availability, legal/regulatory compliance, data residency, performance optimisation, or potentially simply a feel-good factor of having their workloads running locally. Understanding the motivation helps to ensure we provide the right guidance.
- **What's the scope of the workload to be repatriated?** Are they aiming to move all of their Azure resources, those resources related to a specific workload/solution, or just a single resource?
- **Have they considered other alternatives to repatriation?** Other alternatives might include the following:
 - Build a multi-region solution by keeping existing resources and deploying to a new region, or

(where supported) adding a new region to the existing resource. While this might not fit an enterprise workload, it can be a good strategy for SaaS and eCommerce solutions, which need to expand globally.

- Use global routing. Keep the application in its current region, and use Azure Front Door or a [global layer 4 load balancer](#) to route requests across the Microsoft global network for global performance acceleration and security. For more information, see [Network secure ingress implementation](#).
- **How will the migration be sequenced?** They might decide to move everything together, or to move sets of resources together (such as an entire tier of a solution), or to move individual components. If they plan to migrate in stages, will they have resources that need to be accessed from both regions? How will the latency affect the application performance during the period when they run across two regions?
- **Are there components that can't be migrated?** For example, Log Analytics and Azure Backup data can't be moved across regions. If you need to retain logs or backups for a specific duration, the customer might have to keep the resources in the old region until that duration has passed. This might, in turn, introduce operational complexity if those logs or backups need to be accessed.
- **When will they have established landing zone components in New Zealand North?** Consider the [landing zone components](#) that the customer needs to deploy before migrating a workload.
- **Who will do the migration activities?** Will a partner be engaged? Will the customer migrate themselves, or use the partner's skills to execute the migration?
- **Can the migration be rehearsed?**
- **Is the migration a good opportunity to resolve any outstanding technical debt or to make other improvements?** Common examples include:
 - Improvements to monitoring: Can the customer ensure they have monitoring enabled for their important resources, and that their team understands the telemetry as well as how to respond to alerts?
 - Testing: If the customer is building a test environment to validate the migration, can they make the same scripts/processes available for other tests in the future?
 - Automation: If the customer is rebuilding the environment anyway, can they follow best

practices and deploy the infrastructure as code, use modern DevOps practices, and automate their deployment and management as much as practicable?

- Availability zones: Can they make use of availability zones in their solution, if they don't already? Many resources must be configured for availability zones when they're first created, so a migration presents a great opportunity to remediate any configuration that needs to be adjusted.
- DR: Can the customer take lesson from the migration process and apply them to their disaster recovery plan?

Understand how cross-region moves work

There isn't a single approach to moving resources across regions. Some Azure resource types provide built-in support for cross-region moves, while others need to be recreated or migrated by using a tool like Azure Migrate. To execute a move, you'll need to make a plan that considers each resource type that the customer uses. [Azure Region Mover](#) can help to move some types of Azure resources between regions. We provide [guidance for moving many common Azure resource types](#). In general, stateful resources are more complex to move than stateless resources.

Key actions


- **Understand the services and SKUs** that the customer intends to use, and verify they are supported in the region. For more information, see [Service availability and timelines](#).
- **Plan the HA and DR strategy**, and confirm the customer is comfortable architecting with availability zones. For more information, see [Availability zones, zone redundancy, and zonal services](#).
- **Ensure the customer has a support contract** with an appropriate SLA, and that they know how to open support cases. During a migration of a production solution, customers need to be able to quickly reach Azure Support.
- **Verify the permissions in the Azure environment**, including who can create subscriptions, who can set up peerings between virtual networks, and so forth.
- **Plan how the customer's landing zone will be moved to the new region**. For more information, see [Landing zones](#).
- **Plan how connectivity to the new region will work**. For example, the customer might deploy additional ExpressRoute circuits, upgrade an existing circuit, or deploy a new site-to-site VPN. For more information, see [Connectivity to on-premises environments](#).

Considerations

- Consider latency, and whether the application is sensitive to latency. This might influence the strategy and process.
- For many Azure services, DNS names and IP addresses might change when you move across regions.
- Resource names and resource IDs are very likely to change. Verify whether the customer has any automation that depends on finding resources with a specific name or resource ID.
- Azure Reservations are regionally scoped, but [can be exchanged](#) for new reservations in another region.

Resources

- [Cloud Adoption Framework relocation guidance](#)
- [Migrating Azure services to new regions](#)
- [Azure Region Mover](#) can help to move some types of Azure resources between regions.
- We provide [guidance for moving many common Azure resource types](#).



Key technical areas

Service availability and timelines

Key points:

Services will be released to New Zealand North progressively, over a period of several months. Ensure customers understand which services they use, and that they plan based on availability timeframes.

Customers might need to plan workarounds if they need to use services that aren't available immediately upon region launch.

Any customers or partners who plan for a deployment into New Zealand North should be made aware of the service availability and sequencing.

The region will first "go live" for invited customers only. This period is called managed access. After 60 days, the region will be generally available for all customers to use.

Azure services are released to new regions on an ongoing basis. Azure assigns service categories as foundational, mainstream, and strategic. For a list of the services in each category, see [Available services by region types and categories](#). It's important to understand the deployment timelines:


- **All foundational services** will be available from the start of the managed access period.
- **All mainstream services** will be gradually deployed to the region, beginning at the start of managed access and continuing until about 90 days after general availability of the region. We don't know the exact deployment sequence, and we can't prioritise specific services over others. We sometimes have a very approximate idea of sequencing (for example, Private Link often comes early in the mainstream services rollout, and Virtual WAN often comes later) but nothing more specific.

- **Strategic services** are deployed individually based on demand, and we need to provide a business case before deployment can be planned. Strategic services are handled on a case-by-case basis. Please don't make assumptions about whether or when a particular strategic service will be available in the region.

If a customer plans to use New Zealand North soon after launch, it's critical that they evaluate the services and SKUs they intend to use to ensure they will be available. If they won't be available, or if it's unclear, then they should make contingency plans. These plans might include:

- **Wait until their chosen set of services are available** in New Zealand North before they begin to deploy their workload. If they are building a new workload, consider whether they can deploy non-production environments in another region, and then plan to deploy their production services in New Zealand North when their services are ready.
- **Temporarily deploy their workload to another region** with the intention to repatriate it when their chosen services are ready. They might need to accept latency and cross-region traffic cost during the period of time they run across regions. If they follow this approach, [review the guidance for repatriation](#) for some important considerations.
- **Use a different service or SKU** that is available in New Zealand North, and then migrate to the desired service or SKU later.
- **Use a partner service**, such as a network virtual appliance (NVA), while waiting for an Azure-native service to become available.

We acknowledge that there are tradeoffs with each type of contingency plan. Customers and partners need to make an informed decision on the tradeoffs that are right for their scenarios.



Here are some example contingency plans for mainstream services that might not be available immediately on GA:

Service	Example Contingency Plan
Azure Cache for Redis	Customers could deploy their own Redis cache on a virtual machine temporarily, and migrate to Azure Cache for Redis when it's available.
Azure Firewall	Customers could consider deploying a partner firewall as a network virtual appliance (NVA) as an interim step.
	Customers could deploy a firewall into Australia East and use global VNet peering to send traffic to the firewall. Then they can redeploy the firewall in New Zealand North when it's available.
Azure Functions	Customers might deploy their function app to Australia East, and then redeploy to New Zealand North when it's available.
Azure Monitor: Log Analytics	Customers could send their Azure diagnostic and activity logs to a storage account in New Zealand. However, they can't easily query across logs in a storage account.
	Customers can route their logs to a Log Analytics workspace in another region. However, they will pay cross-region traffic.
Azure Virtual WAN	Customers can deploy a hub-and-spoke VNet in New Zealand North, and then migrate to Azure Virtual WAN when it's available.

Please work with your customer's aligned Microsoft seller to ensure that demand for services is recorded in Microsoft's CRM (MSX).

Important: Some of the components required by an enterprise-scale landing zone (ESLZ) are considered mainstream services. This means that a full ESLZ might not be deployable immediately upon region launch. Ensure that customers understand they might need to implement workarounds temporarily.

Strategic services

Not all strategic services will be available in New Zealand North, for a variety of reasons. It's important that customers understand that they might need to run some components of their workloads in other regions if they need specific services that aren't available locally. Alternatively, they might be able to redesign their deployments to use services that are supported in New Zealand North.

Availability zones, zone redundancy, and zonal services

Key points:

Availability zones provide a high degree of redundancy and support resiliency for most workloads' requirements. They are the right choice for most Azure customers.

Availability zones can be used to support a variety of types of workloads and requirements.

Multi-region architectures should be considered when a workload has users who are globally distributed, or when a customer has a mission-critical system that means they are extremely risk averse, and they're prepared to accept the cost and complexity of a multi-region service.

New Zealand North has three independent availability zones. Availability zones are sets of datacentres that are fully isolated from each other. They're physically located far enough apart to reduce likelihood of an issue at one affecting another, but close enough together to allow synchronous replication between zones. For more information about how we design regions with availability zones, see [What are Azure regions and availability zones?](#)

For New Zealand North, there's no paired secondary region. Customers need your help to understand how to use availability zones and other approaches to achieve their resiliency requirements.

Zonal and zone-redundant deployments

There are two ways that resources might be configured to use availability zones. The option you use depends on the way the underlying Azure service works as well as customer requirements.

- *Zonal resources* (also called zone pinning) are deployed into a specific availability zone. This approach doesn't inherently provide any sort of high availability or disaster recovery, but by deploying multiple resources across multiple availability zones you can achieve sophisticated availability requirements. The customer is responsible for handling data replication and failover, but this approach can work well for latency-sensitive workloads. Virtual machines are an example of a zonal resource.
- *Zone-redundant resources* have multiple instances that are deployed across multiple availability zones. Microsoft handles data replication and failover between zones. Many platform as a service (PaaS) services can be deployed in a zone-redundant configuration, such as App Service and Cosmos DB.

In general, we recommend that most production workloads should use multiple availability zones. If possible, use zone-redundant resources, which often give the best set of tradeoffs between resiliency, cost, and operational complexity.

To learn more about how to design a solution with zonal or zone-redundant deployments, and the tradeoffs in the different approaches, see [Recommendations for using availability zones and regions](#).



Disaster recovery with availability zones and regions

A workload's disaster recovery strategy is heavily dependent on the customer's business requirements. It's important that customers have a clear understanding of the business criticality of the system and the implications of downtime so they can make an informed decision about how they should plan for different types of failure.

Availability zones should play a significant role in a customer's disaster recovery strategy. By spreading their workload across multiple availability zones, they mitigate many risks, such as a datacentre outage, power or network connectivity failure, or local weather events. The [Metro DR approach](#) is one way to use availability zones for disaster recovery, with zonal deployments that fail over to an alternative zone. Zone-redundant deployments also provide protection against many of these risks.

Occasionally, customer workloads might be so critical that the customer wants to consider mitigating the risk of an entire region outage. It's important to remember that region-wide outages are extremely unlikely. If a customer needs to mitigate this kind of risk, they need to be prepared to deal with a more complex multi-region architecture. There are complex tradeoffs to consider involving latency, durability, and cost.

For most customers, availability zones provide the best set of tradeoffs to achieve high resiliency without introducing a lot of extra cost and complexity. We recommend most customers use availability zones as part of their primary resiliency approach, and consider cross-region data backups as an extra layer of protection.

Backups

It's usually a good idea for customers to back up their data to another Azure region, even if they are using availability zones as their primary resiliency strategy. In New Zealand, a cross-region backup necessarily means sending data to another geography – most likely Australia. For most customers this isn't a concern. However, if customers have strict data residency requirements that also apply to their backups, they might need to consider exporting their backups out of Azure and to another New Zealand-hosted storage facility.

Explaining availability zones to a customer

When discussing availability zones with customers who don't currently use them, it can be helpful to frame their use as the next iteration of their landing zone approach. This way, it's clear the customer hasn't done anything wrong with their previous architecture, but they and the platform have matured and so they can now take a modern approach to resiliency.

Resources

- [Recommendations for using availability zones and regions](#)
- [Availability zone mappings](#): Each subscription's availability zone order is different. This article describes how subscriptions are mapped to physical availability zones, and how customers can safely use availability zones across subscriptions.
- [Well-Architected Framework reliability pillar, including Design reliable Azure applications](#)
- [Availability zone migration guidance overview for Microsoft Azure products and services](#)

Availability zone-based reference architectures

The Azure Architecture Center is building up a set of reference architectures that use availability zones, including the following:

- [High availability enterprise deployment using App Services Environment](#)
- [IaaS: Web application with relational database](#)
- [Baseline highly available zone-redundant web application](#)
- [Azure Spring Apps baseline architecture](#)



Connectivity to on-premises environments

Key points:

ExpressRoute is a highly available service, with redundancy in each part of the Microsoft environment, including the Auckland PoP. We provide two links for redundancy. Customers need to ensure their equipment and configuration is ready for high availability.


If customers have a mission-critical workload and have low risk tolerance, they can mitigate the risk of PoP outages with a variety of approaches. This is not necessary for most customers.

Customers have multiple choices for how they connect their on-premises environments into Azure, including public internet connectivity, ExpressRoute, VPNs (point-to-site and site-to-site), and SD-WANs.

The Microsoft WAN

Microsoft's wide-area network provides connectivity across our global network of datacenters and edge sites. Each Azure region provides two regional network gateways (RNGs), which are redundant components that enable connectivity between all of the datacenters in the region as well as out to the global Microsoft network.

Customer traffic enters or leaves the Microsoft network at an edge location. We provide two types of edge locations, which are also called points of presence (PoPs). Public edge locations are used for connectivity across the internet, and private peering locations are used to support private connectivity through ExpressRoute.



ExpressRoute

ExpressRoute is often the preferred solution for customers to connect their networks to Azure. Customers can use ExpressRoute today to connect to their Azure virtual networks and resources in Australia and beyond. When New Zealand North is launched, customers can also use ExpressRoute to connect to resources deployed to that region. For a complete list of ExpressRoute connectivity partners in New Zealand, see [Connectivity providers and locations for Azure ExpressRoute](#). Customers should contact their ExpressRoute partner to plan how they will create new circuits or move any existing circuits.

Auckland point of presence

Microsoft provides an ExpressRoute point of presence (PoP), also called a Microsoft Enterprise Edge (MSEE) location, in Auckland. The PoP is generally available now and is located at Vocus in Albany. Nothing about the PoP changes when the New Zealand North region is launched.

When customers have virtual networks in New Zealand North, they'll be able to set up an ExpressRoute circuit in the region and then migrate to the new circuit. For more information, see [How to approach migrating from one ExpressRoute circuit to another ExpressRoute circuit, with a focus on methodical transition of the ExpressRoute Private Peering](#).

Connectivity models

Customers can establish ExpressRoute circuits through a partner or by using ExpressRoute Direct. The Auckland ExpressRoute PoP [supports ExpressRoute Direct circuits](#).

ExpressRoute circuit SKUs

New Zealand and Australia share a geopolitical region, which means that ExpressRoute Standard circuits can connect to resources in either country. Customers can use global peering to connect virtual networks in Australia with those in New Zealand as required, and ExpressRoute can connect across the VNets.

ExpressRoute Local circuits will become available when the New Zealand North region is generally available. The Local circuit type will provide connectivity to the New Zealand North region only.

For more information about ExpressRoute circuit types, see [Azure ExpressRoute pricing](#).

High availability

ExpressRoute is a [highly available service](#), backed by a [99.95% uptime SLA](#). Microsoft's largest customers depend on ExpressRoute to connect their on-premises networks to Microsoft resources.

Microsoft builds high availability into each layer of the ExpressRoute connection. We provide two separate links for each circuit, each terminating in different physical hardware inside the PoP. If the primary link is unavailable the secondary link can continue to be used. Further, ExpressRoute PoPs are designed with a high degree of redundancy and resiliency.

High availability is a shared responsibility. Customers need to use both physical links, and test their configuration and failover regularly. Customers should deploy their ExpressRoute configuration based on our [guidance for high availability](#). Customers must also ensure that their side of the connection doesn't have a single point of failure. Customer premises equipment (CPE) must be evaluated by the customer to ensure that it supports the customer's high availability requirements. At the application tier, workloads must be able to handle retries correctly, because short, intermittent outages of the ExpressRoute connection are normal and expected within the SLA.

Disaster recovery

Outages of an ExpressRoute PoP are extremely rare, and PoPs are specifically designed to ensure they remain operational even in a variety of failure conditions. Occasionally, some customers deploy mission-critical workloads, and have a lower risk tolerance. As part of a broader disaster recovery plan, they might look to mitigate the risk of an entire PoP failing. **For most customers, the cost-benefit ratio of mitigating this risk doesn't make sense**, and they are better to rely on the built-in resiliency of ExpressRoute as described above.

Customers with mission-critical workloads should have explicit, business-driven requirements that indicate the need to mitigate the risk of a PoP outage. These requirements might include a very low recovery time objective (RTO) or recovery point objective (RPO). In these situations, consider a variety of mitigations, each of which have tradeoffs, and remember that in a disaster scenario it's common to run a degraded experience. Mitigations you might consider include:

- Deploy a second ExpressRoute circuit, which will connect to [another PoP in another region](#) (typically Sydney or Melbourne). This option is fully described in the [ExpressRoute disaster recovery guidance](#). Note that this approach will add latency due to cross-region traffic, but in a disaster scenario, added latency is often seen as a reasonable tradeoff.
- Use a [site-to-site VPN as a fallback](#).
- Consider whether a software-defined WAN (SD-WAN) with a partner network virtual appliance (NVA) might provide an alternative connection path.
- Consider other fallback options based on the workload and the customer's requirements, such as:

- Use a virtual desktop (VDI) solution, which might be appropriate for low-latency workloads.
- Deploy point-to-site VPN infrastructure, which users can connect to when their primary connectivity is unavailable. This approach can be more cost-effective than a site-to-site VPN.
- Use a public endpoint and identity-based access controls.

Each mitigation option adds complexity and cost, so it's important to clearly identify whether there's a real need for such a mitigation before commencing. Many of our largest customers rely on a single ExpressRoute circuit.

Calls to action

- **Read our HA and DR guidance.** Specifically, [read the ExpressRoute high availability guidance and disaster recovery guidance](#), and understand the [ExpressRoute SLA](#).
- **Ensure customers understand how to configure high availability for ExpressRoute** including using both physical links and using zone redundant gateways where applicable. Customers should test their configurations regularly.
- **Ensure customers also understand that their equipment (CPE) and application workloads are their responsibility.** If CPE isn't configured for high availability their overall posture is compromised. Similarly, ensure they understand that application workloads that depend on ExpressRoute need to be designed to gracefully retry in the event of occasional connection dropouts.

If a customer indicates they have a mission-critical solution and a low risk tolerance, and that the standard high availability PoP architecture is not sufficient:


- Encourage the customer to provide concrete, business-based requirements, such as an RTO and RPO. Ensure that their architectural decisions are grounded in these requirements and that they aren't overcomplicating their solution when they don't need to.
- Based on their requirements, if the customer needs to mitigate the risk of a complete PoP outage, consider using the mitigations described above after fully evaluating their costs and benefits.

Resources

Customers should review the following patterns for ExpressRoute high availability and resiliency:

- [Connection weighting](#)
- [Autonomous system path pre-pending](#)
- [iBGP preferences](#)
- [iBGP or interior gateway protocols](#)

Customers should always use [bidirectional forwarding detection \(BFD\)](#) to help to prevent lengthy failovers to unpredictable results when circuits flap.



Landing zones

Key points:

Customers with an established landing zone can extend it or move it into New Zealand North when the region is launched and the necessary services are available.

A *landing zone* is a set of resources and configuration that establish the basis for an Azure estate.

It can be helpful to think of Azure landing zones as being like city plans. The architectures of workloads deployed into a landing zone are like plans for buildings in a city. A city's water, gas, electricity, and transport systems all must be in place before buildings can be constructed. Similarly, an Azure landing zone's components, including management groups, policies, subscriptions, and role-based access control (RBAC), all must be in place before any production workloads can be deployed. For more information, see [What is an Azure landing zone?](#)

When a customer deploys a landing zone, many components are not regionally bound. For example, management groups, role assignments, DNS zones, and Microsoft Defender for Cloud configuration are set globally and apply regardless of which Azure regions you use.

There are also some components that are regionally bound, especially for components that relate to networking, connectivity to on-premises environments, Log Analytics and monitoring resources, and automation of management. For more information, see [How landing zones use Azure regions.](#)

Guidance for customers

- If a customer is new to Azure and is deploying a brand-new landing zone after the New Zealand North region launches, they will likely deploy the regional landing zone components directly to New Zealand North.
- If a customer has an existing landing zone and is extending their Azure estate to New Zealand North, many of their existing landing zone resources will remain in place. However, they will likely need to deploy new resources into New Zealand North. For more information, see [Add a new region to an existing landing zone.](#)
- If a customer has an existing landing zone and is migrating their entire Azure estate to New Zealand North, they can move the regionally deployed landing zone components. For more information, see [Move your Azure estate to a new region.](#)
- If a customer has an existing Azure estate but hasn't configured a landing zone, review the guidance at [Brownfield landing zone considerations.](#)
- Important: remember that not all services are available immediately after launch. If a customer needs to deploy early into New Zealand North, they might need to work around service availability for some of the regional landing zone components.

Security considerations for New Zealand North

Key points:

Security for New Zealand North is the same as other Azure regions.

Azure regions all meet the same stringent security requirements, including physical access to datacenters, network security, and hardware- and software-layer security throughout the entire Azure environment. When customers use New Zealand North or any other region, there's nothing different from a security perspective.

Ensure that you understand how services are deployed to New Zealand North and to other regions. Strategic services might not be available in New Zealand North at all. Customers who use specific security products might need to consider running those services in other regions, and will need to determine how this can meet their data residency requirements.



Data residency

Key points:

New Zealand North provides in-country data residency. However, there are some situations where customer data might leave the geography, and customers need to make informed decisions about whether to allow this to happen.

New Zealand-based customers' Microsoft 365 tenancies can be migrated to the New Zealand North region. They must purchase the Microsoft 365 Advanced Data Residency add-on to have their data migrated.

Customers might choose to use New Zealand North to meet data residency requirements for their workloads in the Microsoft cloud.

It's important for customers to be aware that, in some narrow situations, data might be stored outside of their selected geography. For more information, see [Data residency in Azure](#).

Also, if customers are using a wide range of Azure services, they might need to use multiple regions because not all services are available in all regions. They should carefully consider whether their services will be available in New Zealand North. For more information, see [Service availability and timelines](#). If they depend on services that aren't available in the region, they should determine which other regions provide a good balance between their data residency requirements, resource cost, and latency.

Microsoft 365 Advanced Data Residency

Microsoft has recently improved its data residency commitments for Microsoft 365 services with the opening of new regions. This enhancement is particularly beneficial for the public sector and regulated industries, which often have stricter data residency requirements than other customers.

If a customer wishes to migrate from Australia to New Zealand North, they need to purchase the Microsoft 365 Advanced Data Residency SKU in addition to their existing licenses (for example, E3 or E5).

When the region reaches general availability, if the customer has purchased the additional SKU, they can specify the destination for their tenancy migration in the Microsoft 365 portal. This step is crucial to ensure that the customer, as the tenancy owner, has consciously decided to migrate to the chosen destination.

The migration process is managed by Microsoft. Customers can monitor the progress of individual service migrations through their portal. After all services have been successfully migrated, customers will receive a formal notification. This approach ensures transparency and allows customers to stay informed throughout the migration process. For more information see [Advanced data residency in Microsoft 365](#).




Compliance and regulatory standards

Key points:

Microsoft provides New Zealand-specific resources for compliance purposes.

Microsoft Azure and other Microsoft cloud products meet many global and New Zealand compliance and regulatory standards. See [New Zealand regional guidance on the Microsoft Service Trust Portal](#) for compliance and regulatory information.

Also, for the Azure Policy initiative that corresponds to the New Zealand ISM Restricted standard, see [Regulatory Compliance details for NZ ISM Restricted](#).





The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This white paper is for informational purposes only. Microsoft makes no warranties, express or implied, in this document.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in, or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2023 Microsoft Corporation. All rights reserved.

Microsoft and other trademarks are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.