nerdio

THE COMPLETE AZURE-NATIVE DAAS GUIDE:

AZURE HIERARCHY, AZURE VIRTUAL DESKTOP, & WINDOWS 365

Find more:

getnerdio.com

TABLE OF CONTENTS

INTRODUCTION

Over the past two years, interest in cloud-based delivery models for remote work has soared. Microsoft's Azure Virtual Desktop (AVD) and Windows 365 have emerged as ideal Desktop-as-a-Service (DaaS) solutions for securely enabling remote work regardless of an employee's location.

Azure is a diverse cloud platform that contains hundreds of products (also known as SKUs). Azure to cloud is like Apple to devices – each has many SKUs within multiple categories.

The first step in successfully deploying Azure DaaS for your organization is getting to know the two native Azure services, as well as the fundamentals of Azure – its hierarchy, elements, and terminology.

1. AZURE VIRTUAL DESKTOP PLATFORM

The native Azure Virtual Desktop management plane is a Platform-as-a-service (PaaS) offering from Microsoft that removes the need for customers to host their own management infrastructure. Traditionally this would comprise of several brokers, connections, gateways, and licensing servers, which greatly increase the management overhead for the environment.

Within AVD, an organization's administrative responsibility is reduced to the session host servers and their logical assignment to users via workspaces, host pools, and application groups. In combination, these containers offer flexible management and assignment capabilities for the organization's desktop and application estate.

2. WINDOWS 365 CLOUD PC PLATFORM

Whereas Azure Virtual Desktop is a hybrid of PaaS and Infrastructure-as-a-Service (IaaS) components, Windows 365 Cloud PC is a true Desktop-as-a-Service (DaaS) offering. The desktops are dedicated and offered at a fixed price per month. The service comes in two versions: Enterprise and Business. The Windows 365 Enterprise offering also allows Cloud PCs to join existing Azure VNets, offering access to legacy corporate services if required. Windows 365 offers a great experience at a guaranteed price-point per user, which can be helpful for managing budgets.

Read on to understand the most critical Azure elements and how they interrelate with each other.

3. HIERARCHY OF MICROSOFT AZURE

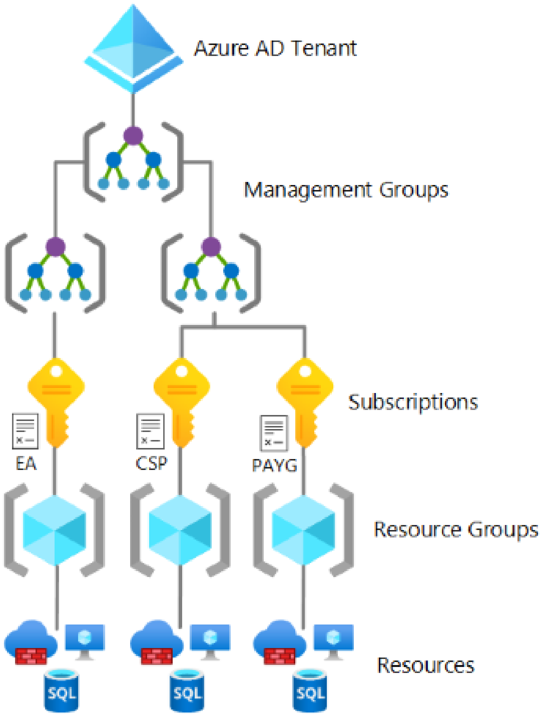

Photo Credit: Microsoft

The image to the right image illustrates the permission flow principles of Azure. As a best practice, administrators should follow Zero Trust principles and the principle of least permissions and only assign access for users at the level it is required. Assigning 'Owner' permissions at the subscription level is a security risk and should therefore be restricted to a small group of senior administrators.

3A. Tenant

Azure services follow a strict hierarchy, and the Azure tenant is at the top of this tree. The tenant tends to be at the organizational level. Most organizations will have a single global tenant that provides the identity and authentication layer for their workforce. Many large enterprises may have multiple tenants serving distinct functions. There may be intentional separation of business areas to accommodate different regulatory standards, or due to the adoption of tenants through mergers or acquisitions. Azure Active Directory (AAD) sits within this tenant and allows admins to cascade permissions to the deeper levels of their environment.

It is important to note that permissions set at the tenant level do not propagate directly to the downstream subscriptions. Management groups can provide a useful bridge here.

NERDIO VALUE-ADD

WHEN DEPLOYING DESKTOPS WITH NERDIO MANAGER FOR ENTERPRISE, ADMINISTRATORS CAN LINK MULTIPLE DISCREET AAD TENANTS WITHIN OUR CONSOLE, ALLOWING FOR THE SIMPLE ADMINISTRATION OF DESKTOPS ACROSS ALL LINKED TENANTS FROM A SINGLE CONSOLE.

4. MANAGEMENT GROUPS

Within the tenant, hierarchical management groups can be created to manage various downstream features of the Azure environment. These management groups can be very useful, especially when managing large, complex environments. Management groups sit above subscriptions, meaning a single management group may be delegated to manage multiple subscriptions.

5. SUBSCRIPTIONS

The subscription serves as the billing container for resources within Azure. A single tenant may contain multiple subscriptions, and these may be procured from various sources (Pay As You Go (PAYG), Enterprise Agreement, or Cloud Service Provider (CSP)). Resources within a subscription are billed in line with the agreement.

While CSP and Enterprise agreements are invoiced according to the agreed terms, it is important to remember here that PAYG subscriptions are charged directly to a credit card, and funds must be available at the time of billing to avoid service interruption. This feature makes PAYG subscriptions unsuitable for critical business services, however they can be valuable for development and testing functions.

Subscriptions were introduced as a discreet billing instance, and still function as a billing container as described above. However, Microsoft now advises that the subscription should also be considered as the top-level management container for a logical collection of resources, as it provides granular access control over all objects it contains. For example, a tenant may contain a 'Finance Team' subscription which contains

NERDIO VALUE-ADD

WITH NERDIO MANAGER FOR ENTERPRISE ADMINISTRATORS CAN LINK MULTIPLE SUBSCRIPTIONS, ALLOWING FOR THE SIMPLE ADMINISTRATION OF DESKTOPS ACROSS ALL LINKED SUBSCRIPTIONS FROM A SINGLE CONSOLE, WITH NO NEED TO SWITCH CREDENTIALS OR CONTEXT.

6. RESOURCE GROUPS

Below the subscription level, we have resource groups (RGs). The RG serves as a bucket for an associated collection of resources. Additionally, specific permissions can be set at the RG level to provide access to it and the resources it contains, without needing to specify these permissions at the subscription level.

RGs contain corporate IaaS and PaaS services and generally group these services in a logical manner. For example, a collection of IaaS virtual machines (VMs) in the West Europe Azure region may sit within a single resource group, along with their virtual disk and Network Interface Card (NIC) objects. This way, an administrator can easily manage the VMs and dependencies for a specified region, and their access can be limited to just the RG that contains these objects. Crucially, resources within a RG do not need to be in the same region as the group. The group is simply a logical management container that allows administrators to manage access to a collection of resources.

Administrators can create policies at the RG level to enforce requirements such as restricting the locations where new resources may be created or mandating the inheritance of tags (object labels).

7. AZURE RESOURCES

At the bottom of the Azure hierarchy tree we have resources. These are the actual objects which constitute corporate services. VMs, SQL databases, firewalls, storage accounts all reside within a RG.
Storage services within the RG may employ geographical resilience, such as Zone or Geo Redundancy (ZRS / GRS). It is important to remember that this is not a disaster recovery solution but rather a data protection mechanism. If a business wants to include disaster recovery for their IaaS services, they should investigate Azure Site Recovery (ASR). Similarly, to provide Highly Available (HA) services, VMs should be created in Availability Zones or Availability Sets, depending on the requirement.

In the following sections we will focus on IaaS VMs and their associated objects but as stated above, it's important to remember that all Azure resources will reside within RGs.

7A. A Word on Tagging Resources

Tagging, or assigning a metadata label to resources, is one of the most useful yet often overlooked features of the Azure ecosystem. All resources including RGs can be assigned multiple tags. These tags can specify the resource's function, version, department, billing allocation or any other label your organization may require. Tags can also be inherited from the RG, meaning they get applied when new resources are created.

Then creating a resource group, it's often useful to specify an 'Owner' tag, with the name or email address of the individual responsible for the resources it contains. This can help prevent resource sprawl, and quickly identify an individual or department who can provide information on the resources it contains.

7B. Virtual Machines

VMs underpin the vast majority of Azure deployments. Many organizations have migrated or are in the process of migrating their legacy on-premises workloads to Azure. Alternatively, some organizations benefit from a hybrid cloud approach where they leverage cloud resources in parallel with their on-premises or hosted infrastructure. While this can introduce some management complexity, it can provide greater flexibility and potentially cost savings, especially for organizations with significant data center or hardware investments.

Azure VMs are offered in a variety of pre-configured sizes or families. These sizes relate to the associated VM resources (RAM, CPU, number of disks) and are all priced differently. There are many different VM sizes available. To help users choose the correct VM size, they are grouped into a set of overarching categories as described below.

TYPE	SIZES	DESCRIPTION
General Purpose	B, Dsv3, Dv3, Dasv4, Dav4, DSv2, Dv2, Av2, DC, DCv2, Dv4, Dsv4, Ddv4, Ddsv4, Dv5, Dsv5, Ddv5, Ddsv5, Dasv5, Dadsv5	Balanced CPU-to-memory ratio. Ideal for testing and development, small to medium databases, and low to medium traffic web servers.
Compute Optimized	F, Fs, Fsv2, FX	High CPU-to-memory ratio. Good for medium traffic web servers, network appliances, batch processes, and application servers.
Memory Optimized	Esv3, Ev3, Easv4, Eav4, Ebdsv5, Ebsv5, Ev4, Esv4, Edv4, Edsv4, Ev5, Esv5, Edv5, Edsv5, Easv5, Eadsv5, Mv2, M, DSv2, Dv2	High memory-to-CPU ratio. Great for relational database servers, medium to large caches, and in-memory analytics.
Storage Optimized	Lsv2, Lsv3, Lasv3	High disk throughput and IO ideal for Big Data, SQL, NoSQL databases, data warehousing and large transactional databases.
GPU	NC, NCv2, NCv3, NCasT4_v3, ND, NDv2, NV, NVv3, NVv4, NDasrA100_v4, NDm_A100_v4	Specialized virtual machines targeted for heavy graphic rendering and video editing, as well as model training and inferencing (ND) with deep learning. Available with single or multiple GPUs.
High Performance Compute	HB, HBv2, HBv3, HC, H	The fastest and most powerful CPU virtual machines with optional high-throughput network interfaces (RDMA).

Source: Microsoft

In its most basic form, a functional Azure VM consists of three dependent objects:

- VM
- Network interface (multiple network interfaces are possible)
- Disk (multiple disks are possible)

These items will typically all reside within a single RG in a single location, as shown below. Ideally, these devices should share a naming convention making them simpler to search for and identify. Tagging can also help to manage associated resources.

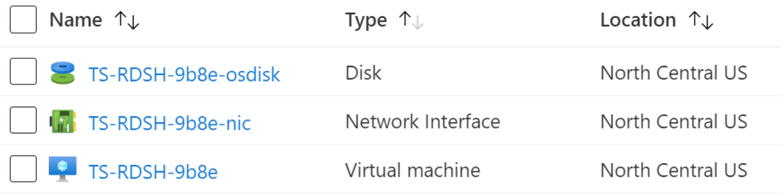

□ Name ↑↓	Type ↑↓	Location ↑↓
□ TS-RDSH-9b8e-osdisk	Disk	North Central US
□ TS-RDSH-9b8e-nic	Network Interface	North Central US
□ TS-RDSH-9b8e	Virtual machine	North Central US

An allocated (running) VM incurs costs per minute, these costs aggregate the items listed above. It's important to remember that a machine that has been shut down from within the OS is still considered to be allocated and will incur costs. To prevent this, the VM must be deallocated by clicking 'Stop' from the console (below), or via a PowerShell command.

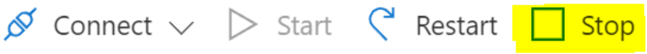

7C. Subscription Core Quotas

Customers are allocated a limited quantity of VM resources by default. This helps Microsoft monitor and manage demand across their global data centers. Administrators can review their per-subscription allocations via the 'Usage + quotas' menu within the subscription (highlighted red in image below).

As shown in the image below, current usage against the quota can be viewed (highlighted blue) and if an increase is required, this can be requested via a simple request process (highlighted green).

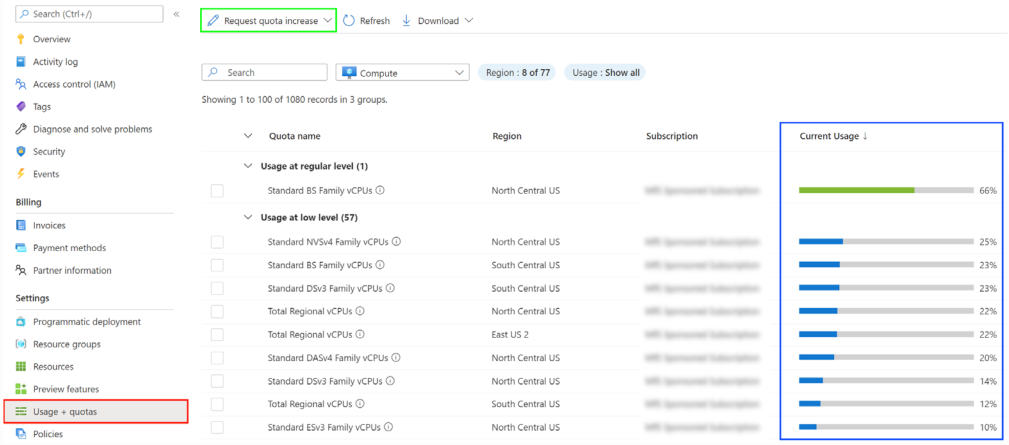

Showing 1 to 100 of 1080 records in 3 groups.

Quota name	Region	Subscription	Current Usage ↓
Usage at regular level (1)			
□ Standard BS Family vCPUs ⓘ	North Central US	[redacted]	66%
Usage at low level (57)			
□ Standard NVSv4 Family vCPUs ⓘ	North Central US	[redacted]	25%
□ Standard BS Family vCPUs ⓘ	South Central US	[redacted]	23%
□ Standard DSv3 Family vCPUs ⓘ	South Central US	[redacted]	23%
□ Total Regional vCPUs ⓘ	North Central US	[redacted]	22%
□ Total Regional vCPUs ⓘ	East US 2	[redacted]	22%
□ Standard DASv4 Family vCPUs ⓘ	North Central US	[redacted]	20%
□ Standard DSv3 Family vCPUs ⓘ	North Central US	[redacted]	14%
□ Total Regional vCPUs ⓘ	South Central US	[redacted]	12%
□ Standard ESv3 Family vCPUs ⓘ	North Central US	[redacted]	10%

Some specialist VM sizes may require additional information including supporting evidence to justify the allocation of resources. Some VM types may not be available in all regions, so if these specific resources are required, it may be that they must be created in a different region. In this instance, administrators should ensure that latency between workloads is tested and confirmed to be acceptable.

7D. Service Level Agreements (SLA)

All Azure services benefit from an SLA, but the targets for these SLAs vary based on the resource type or configuration. For example, VMs within an Availability Zone benefit from a guaranteed 99.99% availability SLA, whereas single-instance machines which are not within an Availability Zone have a guaranteed 99.5% availability SLA.

It is important that organizations design and implement their services appropriately to meet their desired SLA levels within Azure.

7E. Azure Storage

Azure data often resides within Azure storage accounts. Historically, this included VM disks (unmanaged), but the more modern managed disks reside within Azure-managed storage accounts which are not visible to customers.

Storage accounts are generally used to provide serverless file shares or blob storage. The resources within these storage accounts can then be accessed by users or services which have been assigned appropriate permissions.

When creating a storage account, the administrator must decide on whether to create a Standard or Premium account. A Premium storage account provides dedicated file or blob services and provides significantly higher performance than a standard storage account, but at a higher price point. To increase security, storage accounts may be configured to only allow traffic from specified Azure networks, or even to directly communicate with a specified private endpoint to privately connect to a service or resource. This private endpoint must be in the same region as your virtual network, but it can be in a different region from the resource that you are connecting to.

When creating a file share within a Premium storage account, the performance characteristics of the file share are a function of its provisioned size. As shown below, this increases the file share size dynamically and increases the available IOPS and throughput rate.

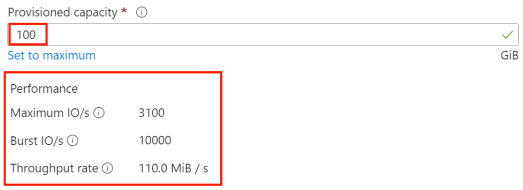

Provisioned capacity * ⓘ
100 GiB

Performance
Maximum IO/s ⓘ 3100
Burst IO/s ⓘ 10000
Throughput rate ⓘ 110.0 MiB / s

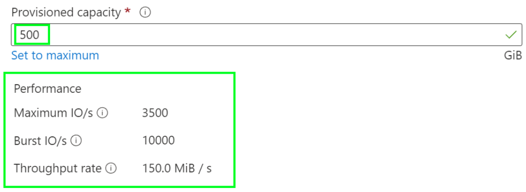

Provisioned capacity * ⓘ
500 GiB

Performance
Maximum IO/s ⓘ 3500
Burst IO/s ⓘ 10000
Throughput rate ⓘ 150.0 MiB / s

If the performance requirement goes beyond what is possible with a Premium storage account, then the administrator may choose to investigate the use of an Azure NetApp File Share to provide the required level of performance. Azure NetApp Files require access to an Azure network (Azure vNET).

As is the case with all other Azure IaaS resources, storage accounts reside within RGs. Soft delete, backup and redundancy options are available, the latter of which is discussed in the following section. Some functions are not supported across all data scenarios; therefore, organizations should ensure their approach takes into account the available data protection options.

7F. Data Redundancy

As touched on earlier in this document, Azure Storage accounts can be created with various redundancy options. Data in Azure is replicated three times in the source (primary) region. Administrators may choose how this replication is achieved.

Locally Redundant Storage (LRS) copies data synchronously three times within a single physical location in the primary region. LRS is the least expensive replication option but isn't recommended for applications requiring high availability or durability.

Zone-Redundant Storage (ZRS) copies data synchronously across three Azure availability zones in the primary region. For applications requiring high availability, Microsoft recommends using ZRS in the primary region and replicating to a secondary region.

Additionally, Azure offers the option for data to be replicated to a secondary region. This region is pre-determined by the chosen primary region and cannot be edited by the administrator.

According to Microsoft's website, Azure Storage offers two options for copying your data to a secondary region:

Geo-redundant storage (GRS) copies your data synchronously three times within a single physical location in the primary region using LRS. It then copies your data asynchronously to a single physical location in the secondary region. Within the secondary region, your data is copied synchronously three times using LRS.

Geo-zone-redundant storage (GZRS) copies your data synchronously across three Azure availability zones in the primary region using ZRS. It then copies your data asynchronously to a single physical location in the secondary region. Within the secondary region, your data is copied synchronously three times using LRS.

7G. Azure Networks

Azure allows for the creation of self-contained virtual networks. Azure vNets are comprised of one or more CIDR address spaces, one or more subnets within the specified address spaces, and one or more DNS server addresses. By default, Azure specifies its own DNS servers. These should be replaced by an 'on-network' or domain DNS server if internal VM DNS resolution is required.

These networks can be peered (connected) directly to other Azure vNets, with firewalls or Network Security Groups (NSGs) attached to control traffic. Routing is handled by Microsoft, but route tables can be created if required.

While all the information for a Azure vNet's configuration can be found in the console, as with most Azure resources this can be queried more efficiently via PowerShell. Azure provides a web based 'cloud shell' in the Azure web portal, allowing for simple interrogation of the environment.

The below command will report on the full configuration of a specified Azure vNet.

```
GET-AZVIRTUALNETWORK -NAME <VNET NAME>
```

The output of this command will look similar to the image below. We can see that this Azure vNet has a single /16 address space, with a single /24 subnet and no Azure vNet peering configured.

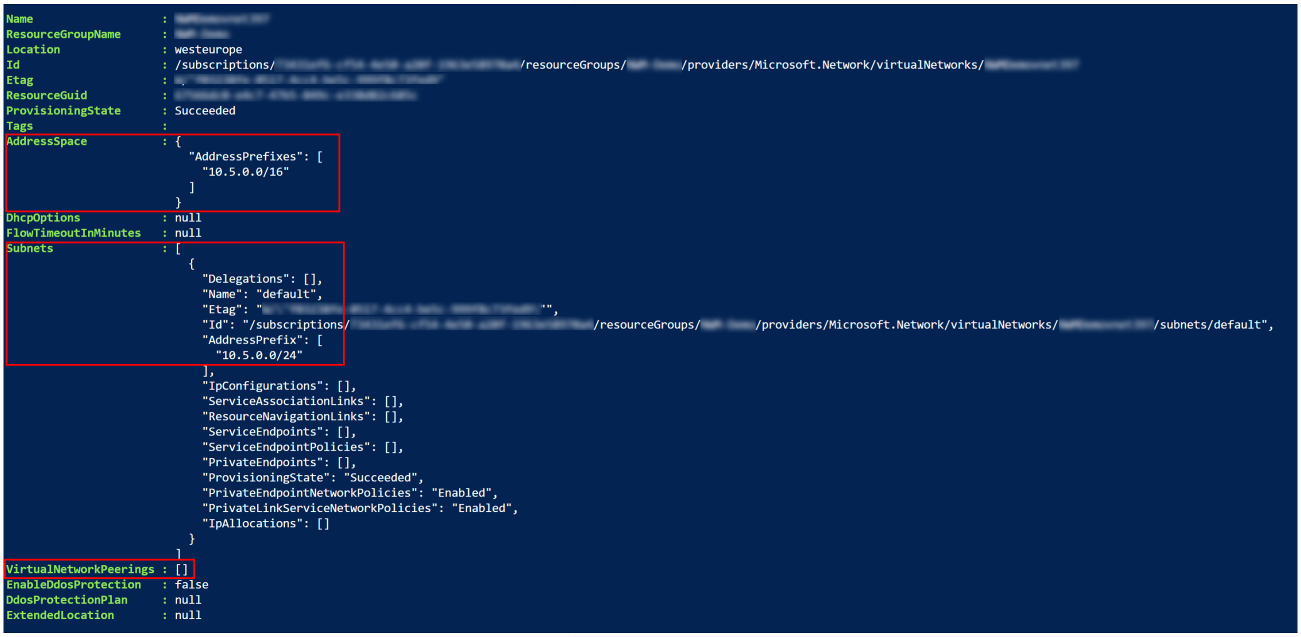

```
Name                   : [redacted]
ResourceGroupName      : [redacted]
Location               : westeurope
Id                     : /subscriptions/[redacted]/resourceGroups/[redacted]/providers/Microsoft.Network/virtualNetworks/[redacted]
Etag                   : [redacted]
ResourceGuid           : [redacted]
ProvisioningState      : Succeeded
Tags                   :
AddressSpace           : {
                           "AddressPrefixes": [
                             "10.5.0.0/16"
                           ]
                         }
DhcpOptions            : null
FlowTimeoutInMinutes   : null
Subnets                : [
                           {
                             "Delegations": [],
                             "Name": "default",
                             "Etag": "[redacted]",
                             "Id": "/subscriptions/[redacted]/resourceGroups/[redacted]/providers/Microsoft.Network/virtualNetworks/[redacted]/subnets/default",
                             "AddressPrefix": [
                               "10.5.0.0/24"
                             ],
                             "IpConfigurations": [],
                             "ServiceAssociationLinks": [],
                             "ResourceNavigationLinks": [],
                             "ServiceEndpoints": [],
                             "ServiceEndpointPolicies": [],
                             "PrivateEndpoints": [],
                             "ProvisioningState": "Succeeded",
                             "PrivateEndpointNetworkPolicies": "Enabled",
                             "PrivateLinkServiceNetworkPolicies": "Enabled",
                             "IpAllocations": []
                           }
                         ]
VirtualNetworkPeerings : []
EnableDdosProtection   : false
DdosProtectionPlan     : null
ExtendedLocation       : null
```

To extend connectivity to a datacenter, on-premises, or multi-cloud environment, administrators have two options-- a site-to-site VPN (other VPN configurations are also supported) or a Microsoft ExpressRoute connection.

7H. Site-to-Site VPN

A site-to-site Virtual Private Network (VPN) is configured by creating a dedicated Gateway Subnet and gateway with an associated public IP address within a nominated Azure vNet. The administrator must also create a VPN connection on their local firewall appliance and assign a public IP address to this. These two endpoints are then configured to create an encrypted connection. As best practice, dual redundant VPN connections should be created to provide resilience and mitigate the loss of a single connection.

7I. ExpressRoute

An ExpressRoute is a Microsoft-certified connection providing built-in resilience through the provision of two diverse-path fiber connections. Due to these strict requirements, ExpressRoute circuits may only be provisioned between ExpressRoute enabled datacenters and the Azure cloud. ExpressRoutes is a private, unencrypted connection between the datacenter and Azure. If traffic encryption is desired, an S2S VPN can be provisioned within the ExpressRoute. This will incur a small performance penalty due to the encryption overhead, as is the case with all encrypted VPN connections.

ExpressRoutes also require a gateway connection to work; however, the details for this connection will be provided by the datacenter service provider.

7J. Data Transfer

Data transfer into Azure (ingress) is free. However, if a large volume of data must be transferred, the bandwidth between the site and Azure may make transfer times unacceptable. In this instance, Microsoft offers the Databox and Databox Heavy data import service.

Data transfer within the same availability zone is free, however data transfer between zones/ regions or outside of Azure is chargeable. ExpressRoutes provide fixed-cost bandwidth between a datacenter and Azure, therefore no additional data ingress or egress charges are payable.

nerdio

Nerdio empowers IT Professionals to deploy, manage, and auto-scale virtual desktops in Microsoft Azure. Created to address the technical and security requirements of enterprise customers, Nerdio Manager for Enterprise is ideal for IT Professionals looking to deploy and manage large Azure Virtual Desktop (AVD) or Windows 365 environments. The platform can be connected to an existing AVD setup in minutes or used to stand up a new AVD deployment in hours. Nerdio Manager is an all-PaaS Azure application that runs in a customer's own Azure subscription, making it one of the most secure and compliant solutions on the market. For more information visit www.getnerdio.com

CONTACT US

Email: hello@getnerdio.com

Website: getnerdio.com

If you would like to learn more about Nerdio and how we can help optimize your Azure environment and Azure DaaS strategy, schedule some time to talk with one of our technical sales experts and see the platform in action.

LEARN MORE

SCHEDULE A DEMO