Findings and Recommendations

Cloud Assessment and Design

Prepared for:

ABC Company

Prepared by:

Neal Zimmerman
Senior Cloud Architect
nealz@myITTeam.pro

February 31, 2025

Table of Contents

Business Objectives

ABC Company environment needs a review to determine cloud opportunities across its existing infrastructure in Montvale, New Jersey and Toronto, Canada.

Microsoft Azure Infrastructure consists of various technologies that together expand the capabilities and value that customers can realize from a cloud. My IT Team is designed to help organizations build hybrid and cloud infrastructures for providing proactive solutions to your top infrastructure challenges. We do this by enabling resilient environments with automation, self-service, and security that customers require to deploy business-critical applications that respond to growing business demands.

This document provides both logical and physical design recommendations encompassing components that are pertinent to this requirement. To facilitate the requirements of ABC Company, these considerations and recommendations are based on a combination of Azure best practices and specific business requirements. Cloud infrastructure–related components, including requirements and specifications for virtual machines, security, networking, storage, and management, are included in this document.

Executive Summary

This architecture is developed to support consolidation and migration of 89 existing on-premises virtual machines to Microsoft Azure Cloud. The required infrastructure defined here can be used as a foundation for any implementation projects to migrate the virtual machines to Microsoft Azure. Microsoft Azure is being adopted to decrease total cost of ownership, reduce the need for expensive datacenter expansions, increase operational efficiency and capitalize on higher availability with elasticity that comes in Azure workloads.

This reference architecture is designed to provide scaling, optimization, and automation when implemented in Microsoft Azure Cloud.

Requirements and Assumptions

The primary requirement for this architecture is to lower the costs, increase agility and elasticity while operational effort involved with deploying workloads should be decreased.

Throughout this design document, adherence to the standards and best practices as defined by Microsoft Azure are recommended when and where aligned with the requirements and constraints as listed in the following sections.

Requirements

ID	REQUIREMENT
R001	Business agility and flexibility should be increased; the cost of doing business should be decreased.
R002	Availability of services is defined as 99.95 percent during core business hours.
R003	Security compliance requires network isolation for specific workloads from other services.
R004	Minimal workload deployment time.
R005	A separate management VNET must be used for shared services.
R006	The environment should be scalable to enable future expansion.
R007	Resources should be guaranteed for groups of workloads as part of internal SLAs.
R008	The recovery-time objective in the case of a datastore failure should be less than 8 hours.
R009	Resiliency should be factored in.

Table 1 - Requirements

Assumptions

ID	ASSUMPTION
A001	Internet connectivity is already provided at primary sites.
A002	Availability of services required to be 99.95 percent during core business hours.
A003	Previous architecture is 3 tiers with Web, Business and Data Layer
A004	All the servers are assumed to be Windows and Linux with Database as MSSQL.
A005	All the other requirements can be incorporated in subsequent phases.
A006	There are standard Microsoft Azure constraints on certain services.
A007	On-premises Infrastructure configuration needs to be provisioned for any dependent service.

Table 2 - Assumptions

PaaS Overview

Azure offers a Platform as a Service (PaaS) solution. PaaS options are easier to configure and include licensing for Operating Systems, Applications & Databases.

Some examples where PaaS can provide options:

Instead of running…	Consider using…
Active Directory	Azure Active Directory
IIS	App Service
NoSQL DB	Cosmos DB
SQL Server	Azure SQL Database
File share	Azure NetApp Files

Table 3 - PaaS Options

The managed instance deployment option preserves all PaaS capabilities (automatic patching and version updates, automated backups, high availability). You may be able to bring your own license through Software Assurance called the Azure Hybrid benefit.

Benefits of PaaS:

- Reduces administrative costs as maintenance, backups, HA, and patching is handled by the cloud provider
- Eliminates hardware costs

Cons of PaaS:

- Higher monthly cost than IaaS
- Limitations to the features available

IaaS Overview

Azure offers an Infrastructure as a Service (IaaS) solution and provides virtual machine instances with compute and storage. These are self-managed VMs. This offers full flexibility in what you can do, but also requires you to do all patching, backups, and other maintenance. You may be able to bring your own license through Software Assurance called the Azure Hybrid benefit.

Azure's virtual machines are implemented as either a single VMs or VM Scale sets (Pooled virtual machines) with access to shared resources.

Benefits of IaaS:

- Eliminates hardware costs
- Access to all features of SQL based on your edition and version
- Lower monthly cost than PaaS

Cons of IaaS:

- All administrative tasks are your responsibility

Scope

Cloud Strategy

The following approach will be used for this Architecture document:

- The document will use guidance on architecting solutions on Azure using established patterns and practices from the Azure Architecture Center.
- The migration strategy proposed in this document will be iterative using Microsoft Cloud Adoption Framework for Azure (See figure below).
- The initial focus of this strategy will be greenfield with a Minimal Viable Product (MVP) approach for standing up workloads within Microsoft Azure.
- Since Microsoft has extensive documentation for Azure, this document will provide a concise summary for each in-scope service or component and a link to the relevant MS documentation.

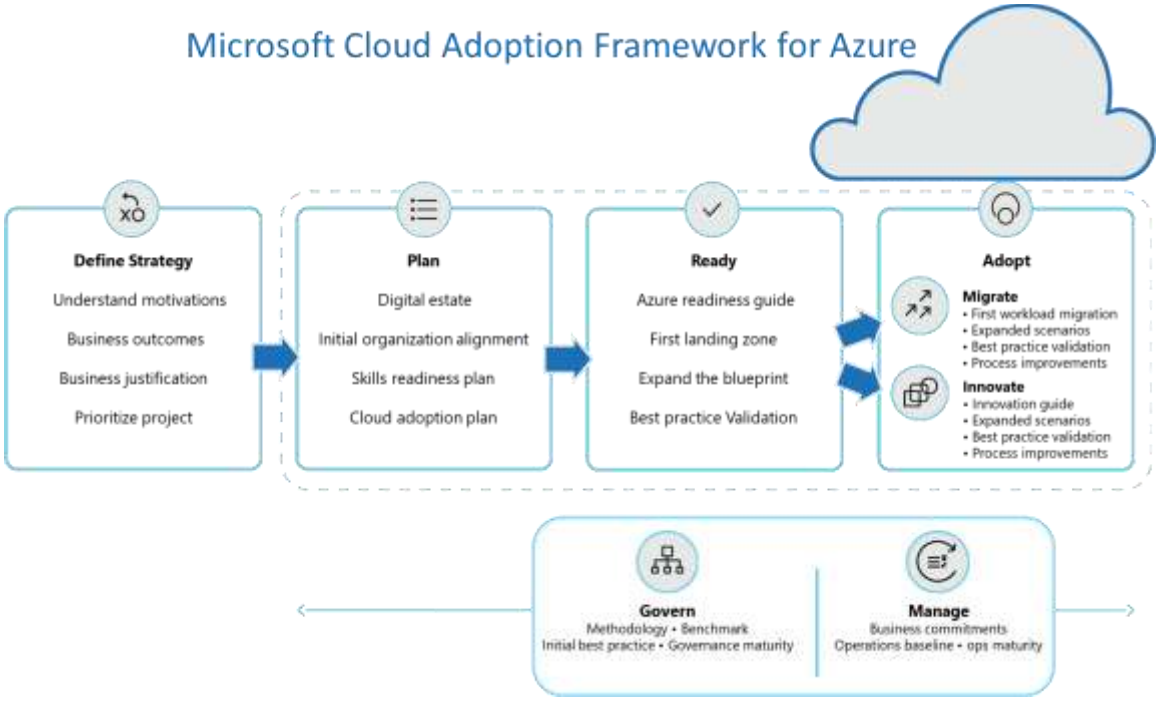

Figure 1 - CAF

Cloud Recommendations

Cloud recommendations are based on overall performance of systems, existing features in use and the ability to shift those to the cloud.

The existing local network will need to be extended into a secure network in Microsoft Azure. This architecture document recommends a secure network, between the on-premises network and an Azure virtual network. This VNET configuration will allow you to take advantage of cloud IaaS or PaaS in conjunction with existing infrastructure. A VNET also gives you control over DNS, traffic between subnets, security, and accessibility.

Implementation of the Azure Virtual Network is outside the scope of this assessment, but it is important to note it is a prerequisite for moving workloads to the cloud.

Important: All current cloud-based offerings offer compute power in the form of vCPUs. One vCPU is the equivalent of a single thread of a hyperthreaded CPU core. While providers claim these perform as well as true cores, real-world experiences indicate that vCPUs perform at roughly 75% of a comparable core.

When moving to the cloud, we recommend allocating 2x the compute power required to bridge the performance gap between vCPUs and true cores. For example, if server A needs 9Ghz of compute power currently, it would be allocated 18Ghz worth of vCPU power.

Resource Requirements

Resource requirements and recommendations are based on the performance data gathered over the course of one week. When resource usage is at 100%, a best estimate is made of the required resources to increase the resource without over-allocating. When resource usage is less than 100%, a conservative estimate is used when recommending lesser resources.

In Design Scope

To meet the Minimal Viable Product requirements and timelines, Azure design will be based on Azure Native services only to support the Greenfield deployment for ABC Company infrastructure and the associated applications.

The Azure cloud platform/foundation design will cover:

- Greenfield Design is based on Azure Native Services that have a scalable foundation to serve ABC Company and other businesses with related service capabilities and should operate independently to other customer organizations that may already be in place.
- The Azure ABC Company Cloud platform will be designed to scale to deliver new cloud services for additional items that ABC Company users may require for future projects.
- Design on using Azure native services.
- Governance Model
    - Enrolment Model
    - Subscription Model & Management
    - Naming Standards
    - Azure Policy
    - Auditing
    - Resources Groups
    - Tagging Standards
    - Role-Based Access Controls (“RBAC”)
    - Resource Locks
- Platform Architecture
    - Hybrid Cloud Network Design
        - Hub & Spoke Reference Architecture
        - VNET Topology
        - On-premises connectivity
        - Network Deployment and Integration with customer WLAN to Private Cloud
        - Secure connection via Virtual Network Gateway from Customer WLAN and segregated LAN (“LAN”) for Other Services (If Any)
        - Windows 10 Point to Site (“P2S”) VPN to Azure
        - IPSEC Site to Site (“S2S”) VPN to Azure
- Security
    - Identify and Access Management (“Cloud Only Identity”)
    - Azure Information Protection
- Compute

    - Storage
- Backup and Restore
    - Disaster Recovery (Resilient Service Design)
- Monitoring and Management
- Automation

Out of Scope

The overview and design exclude:

- Procurement of any tools, hardware and software licenses including Azure Cloud Subscription
- Implementation of any 3rd party tools/software has been excluded from the minimum viable product ("MVP") scope of work, depending on the detailed design any 3rd party tool can be implemented in next phase of the project
- Build and remediation of any existing infrastructure services
- Build of End user management tools such as SCCM or Intune
- Build using DevOps platform for automation
- Patch Management for End User Devices
- Existing customer reference architectures and customer practices
- Federation with any domains
- Classification of Data based on other customer entity setups
- Modern Token / Key Based Authentication for Application-to-Application integration

Terminology and Acronyms

The following table lists the common terms and acronyms that are used throughout the document.

Term/Acronym	Definition
AD	Active Directory
AD DS	Active Directory Domain Services
API	Application Programming Interface
ADE	Azure Disk Encryption
ARM	Azure Resource Manager
ASG	Application Security Group
ASR	Azure Site Recovery
AV	Antivirus
DB	Database
DEV	Development
DMZ	Demilitarized Zone
DR	Disaster Recovery
DSC	Desired State Configuration
EPTM	Endpoint Threat Management
GRS	Geo-Redundant Storage
HA	High Availability
HDD	Hard Disk Drive
HSM	Hardware
IaaS	Infrastructure as a Service
IAM	Identity and Access Management
IOA	Indicator of Attack
IOC	Indicator of Compromise
IOPS	Input / Output Operations Per Seconds
IPSec	Internet Protocol Security
IPS	Intrusion Prevention System
JSON	Java Script Object Notation
LAN	Local Area Network
LRS	Locally Redundant Storage
MFA	Multi-Factor Authentication
MVP	Minimum Viable Product
NGFW	Next Generation Firewall
NIC	Network Interface Card
NSG	Network Security Group
NVA	Network Virtual Appliance
ONR	Office for Nuclear Regulations
OMS	Operations Management Suite
OS	Operating System
P2S	Point-to-Site
PaaS	Platform as a Service
RA-GRS	Read-Access Geo-Redundant Stores
RBAC	Role Based Access Control

SaaS	Software as a Service
SIEM	Security Information and Event Management
SLA	Service Level Agreement
S2S	Site to Site
SSD	Solid State Drive
SSE	Storage Service Encryption
SSL	Secure Socket Layer
SQL	Structured Query Language
TLS	Transport Layer Security
UDR	User Defined Routes
UAT	User Acceptance Testing
VHD	Virtual Hard Disk
VM	Virtual Machine
VNet	Virtual Network
VPN	Virtual Private Network
WSUS	Windows Server Update Services
ZRS	Zone-Redundant Storage

Table 4 - Terminology

Azure Governance Foundations

Adopting the cloud is a journey, not a destination. Along the way, there are clear milestones and tangible business benefits. The final state of cloud adoption is unknown when a company begins the journey.

Cloud governance creates guardrails that keep the company on a safe path throughout the journey. However, without appropriate guardrails in place, businesses can soon find their Azure environment and subscription spend out of control.

It is important to define an Azure governance model that addresses the above concerns but does not overly hinder the agility and flexibility that make Microsoft Azure so attractive to the business. We shall help to define the Azure governance model referencing Microsoft's Cloud Adoption Framework governance model.

The figure below shows the components of the governance model. Its foundation relies on corporate policies that drive governance. The pillars support the corporate policies to avoid potential pitfalls. Microsoft uses what it calls hierarchical "scaffolding" to help organizations build flexible controls over Azure governance policies and accommodate a variety of organizational needs.

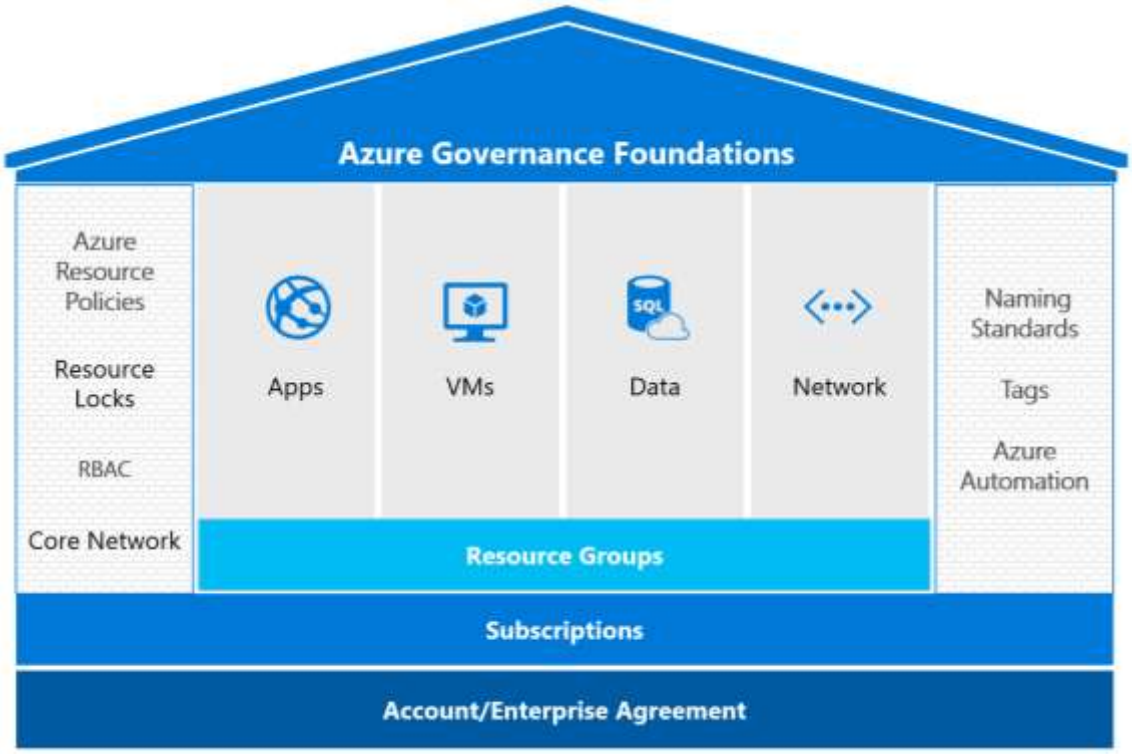

Figure 2 - Azure Scaffold Model

Azure Subscription Hierarchy

The Enrollment model or Subscription hierarchy is key to a successful governance model within Azure. Selecting the most suitable model is a major decision and must be based on the current business requirements but also with an eye to the future.

Enterprise Enrollment - An enrollment is the control point for all Azure services and usage within an enterprise. Admins typically use enrollment to consolidate billing and allocate costs to various business units, projects, and workgroups. Enterprise Agreement customers receive an Azure enrollment number and access key when they initially sign up for Azure. Most enterprises will have one enrollment.

Departments -These provide a means to subdivide Azure resource privileges, usage, and billing within a large organization. Departments are optional, but they help partition many Azure resources into logical units that correspond to a business group, development project, application, or any other organizational structure.

Azure department administrators have management authority over groups of accounts and Azure subscriptions.

Accounts - These are more granular Azure resource and usage controls that admins use for reporting and to manage access to underlying Azure services. The account creator is the default account administrator and controls all Azure subscriptions and the services available to them within an account.

This makes accounts, primarily, a billing construct. Before individual developers begin to use Azure, they must create an account that is tied to a unique ID and credit card number.

Subscriptions - Subscriptions are the level where users create and consume Azure resources. A subscription can also help an organization enforce limits; for example, a limit could prevent the accidental deployment of a massive number of resources, such as VMs, that results in high monthly costs.

This makes subscriptions provide a mechanism to control the Azure services available to individual users and workgroups and to create three parameters: a unique subscriber ID, a billing location, and a group of available resources.

We are proposing following Enrollment model.

1. Enterprise Enrollment (Thru CSP)
2. Departments
3. Accounts
4. Subscriptions

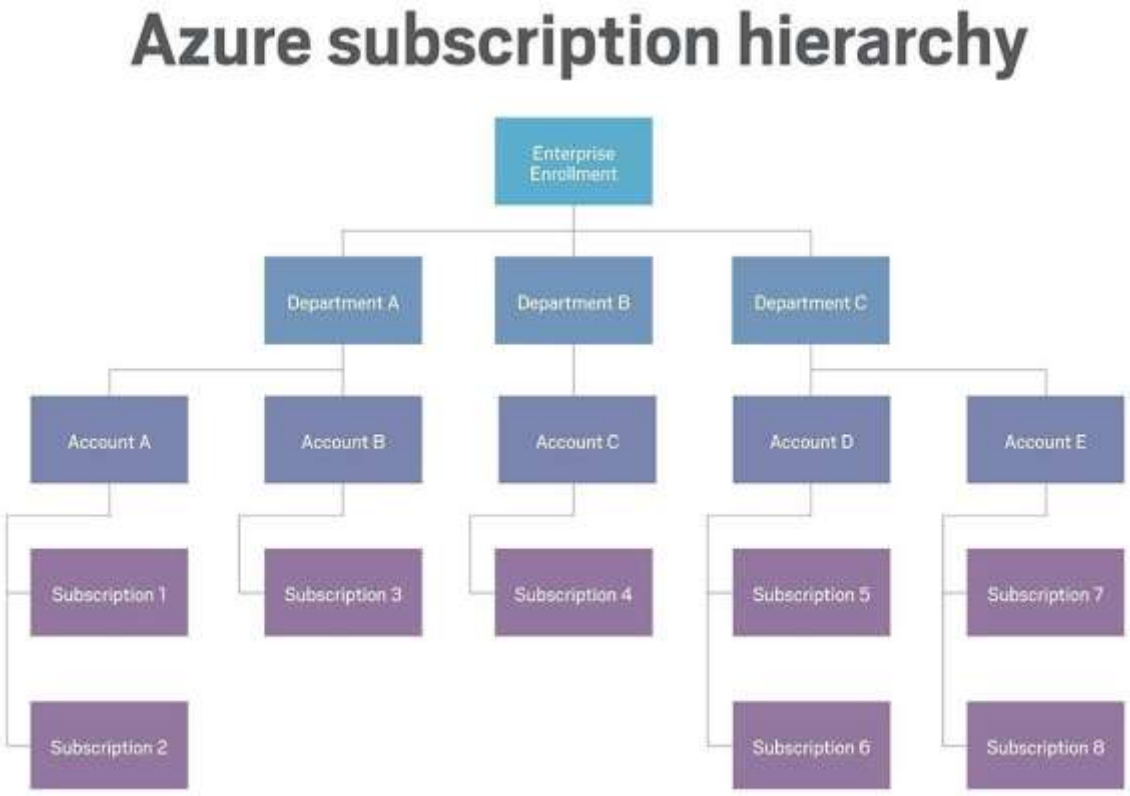

Figure 3 - Azure Subscriptions

Management Groups

Microsoft has recently released new way of modelling an Azure Hierarchy called Azure Management Groups. Management groups are much more flexible than departments and accounts and can be nested up to six levels. Management groups allow you to create a hierarchy that is separate from your billing hierarchy, solely for efficient management of resources.

Naming Standards

Well defined naming standards ensure consistency throughout subscriptions and resources, ensuring resources are easily identifiable within the Azure Portal, in bills and within scripts.

Microsoft Azure Standard Naming Conventions should be followed.

Resource Groups

With Azure Resource Manager, resources can be grouped into meaningful groups for management, billing, or their lifecycle.

Note
Resource Groups cannot be nested/contained within each other, and resources can only belong to one resource group.

It is important to standardize how Resource Groups (RG) will be used throughout the organization. Two common approaches to Resources Groups are:

Lifecycle - Grouping of an entire application allows individual application management.

Management – Grouping of resources based on their role, e.g. All SQL Servers, all Web Servers, all Middleware servers. This approach would allow different teams to manage their respective resources.

Recommendation
We will create a minimum of 4 resource groups: • **Hub-RG** • **Prod-Spoke-RG** • **UAT-Spoke-RG** • **DEV-Spoke-RG**

Tagging Standards

All Resources within Azure can be assigned tags, labelling them with specific information. The tags can be used to aggregate and group resources for reporting, billing or simply be used to provide more details about a resource.

A solid tagging standard provides the metadata required for business, finance, security, risk management, and overall management of the environment.

Recommendation
The following Resource Tags should be used at a minimum: • **Environment** • **Application**

Recommendation
• **Role** • **Tier** • **Department (or Business Unit)** • **Application Owner** • **Cost Management**

Azure Policy

Azure Policy provides the ability to manage risk in Azure; it is a core component of good IT governance within Azure. Policies can be defined and implemented to restrict, enforce, or audit certain actions within Azure, ensuring the standards defined. Examples of how Azure Policy would be used are:

- Geo-compliance
- Cost Management
- Controlling types of resources that can be provisioned
- Tagging Governance

The first step in configuring policies is defining them. A policy definition consists of the condition under which it is enforced and the effect that it has once those conditions are met.

Azure Policy has some of the following built-in policies.

Allowed Storage Account SKUs - This policy definition has a set of conditions/rules that determine if a storage account that is being deployed is within a set of SKU sizes. Its effect is to deny all storage accounts that do not adhere to the set of defined SKU sizes.

Allowed Resource Type - This policy definition has a set of conditions/rules to specify the resource types that your organization can deploy. Its effect is to deny all resources that are not part of this defined list.

Allowed Locations - This policy enables you to restrict the locations that your organization can specify when deploying resources. Its effect is used to enforce your geo-compliance requirements.

Allowed Virtual Machine SKUs - This policy enables you to specify a set of virtual machine SKUs that your organization can deploy.

Apply tag and its default value - This policy applies to a required tag and its default value if it is not specified by the user.

Enforce tag and its value - This policy enforces a required tag and its value to a resource.

Not allowed resource types - This policy enables you to specify the resource types that your organization cannot deploy.

Cost Management - Cost Management is involved in controlling cost and usage of cloud resources with the goal of creating and maintaining a planned cost cycle.

Resource Locks

Resource Locks restrict the modification (Read-only) or deletion (Can Not Delete) of resources, even if the user has the required RBAC permissions for the type of resource they are making the change to. Locks can be applied at a subscription, resource group or resource level.

To modify/delete a resource with a Resource Lock applied, the lock must first be removed. Of the built in RBAC roles, only Owner can create or delete locks.

Given the restrictions and overhead Resource Locks incur, it is recommended they are only used on high value resources that would cause major disruptions to the business if changed/deleted.

One example of such a resource would be the core networking resources for the Azure environment.

Recommendation
Resources Locks should be created for Core Resources, preventing them from being changed and cannot be deleted.

Conceptual Architecture

The document aims to reduce operational overhead and TCO by simplifying management tasks and abstracting complex processes. Throughout this architecture design, all components will be described in depth, including design considerations for all components. The focus of this architecture is resource segregation and isolation.

The environment has the following major pillars:

- Compute
- Networking
- Storage
- Security

Each of the pillars will be carved into multiple pools to provide different service levels for the various workload types. The requirement is to provide a secure and shielded environment.

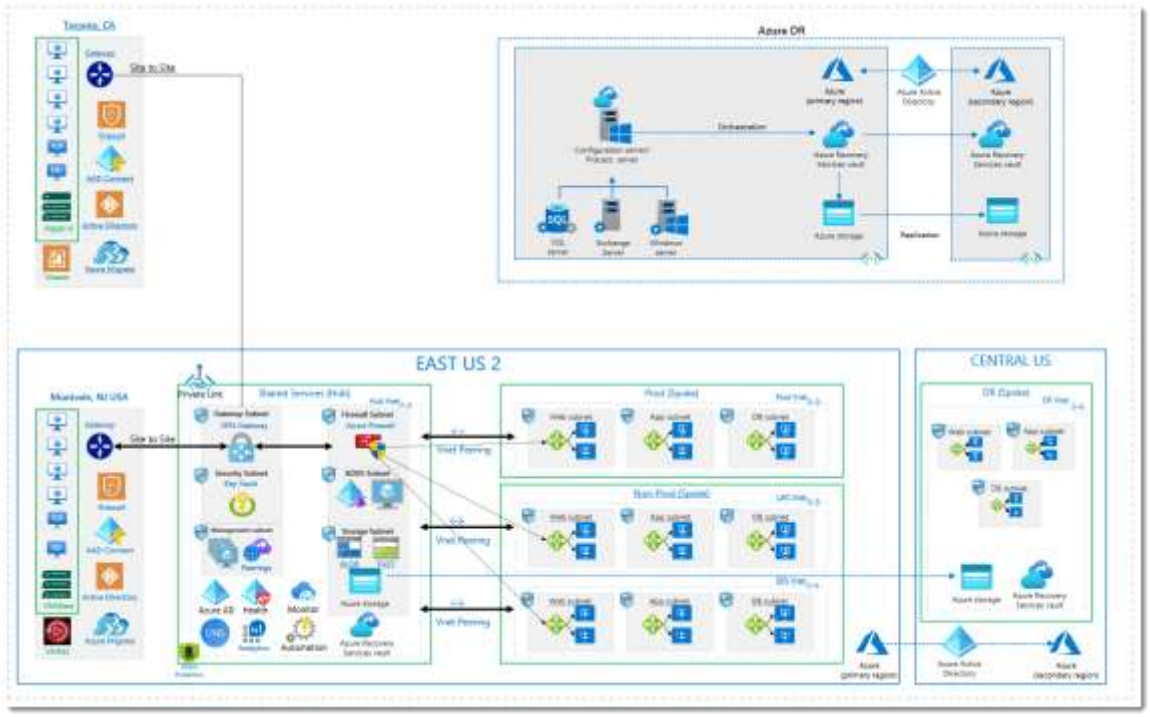

Figure 4 - Reference Architecture

Reference Architecture attached → hybrid-network-hub-spoke-PM-v1.vsdx hybrid-network-hub-spoke-PM-v1.pdf

To meet these requirements, Azure DDoS protection will be implemented to enable the use of security policies on a VNET level. Administrators can define and enforce granular policies for all traffic that crosses a virtual network, increasing visibility over internal traffic while helping to eliminate detours to Azure firewalls. Additional blocks might be added based on requirements and the different types of workloads being deployed.

Platform Architecture

The platform to be created will have traffic coming from

- On premises
- Internet (External Connections)

To handle the traffic the following solutions will be deployed

- The platform will have multiple VNETs created to isolate traffic and allow environment segregation
- IPSEC VPN connection will be established between on premise and Azure US East and US West clouds for interactions
- S2S VPN will be established as redundant connection to Azure US East (Primary) and US West (Secondary) route as fall back option
- Azure default provides DDoS protection at its layer for the incoming traffic.
- Azure firewall will be deployed in the HUB VNet through which all the application traffic will be routed
- Azure User Defined Routing ("UDR") will be deployed for the routing and NSG will be utilized for subnet level communications
- Managed disks to be used for all VMs
- Encryption at transit is (Key Vault) enabled by https connections and rest is enabled with storage account encryption and Azure SSE
    - Option 1: Microsoft-managed keys (Encryption at REST)
    - Option 2: Shared Access Signature
        - Window 10 client to have a mapping utilizing Azure Storage File Share.
- Folder Permissions to be managed by Azure Active Directory native groups.
- Backup will be handled using Azure backup, NSG will be configured in the spoke VNETS to allow communication to azure master servers for enabling backups.
- Monitoring will be handled using Azure monitor for infrastructure.

- Network monitoring will be handled using Azure Network Watcher for the traffic happening within the platform.
- Azure storage account and ARM templates will be utilized wherever possible.

The details of each of these solutions proposed are elaborated in the subsequent sections.

Network

The Network architecture within Azure provides integration to on-premises and takes careful consideration to services that will be consumed at present as well as in the future. The design should allow minimal disruption to a working environment.

Virtual Peering, Zones and Resilience with VPN and Firewalls have been considered. The customer is enabled on-premises on remotely to access services within the Azure cloud.

Each Virtual Network (VNet) specified has gone through careful consideration when specifying the address space as it is understood. The IP range for the entire network should be divided into appropriate subnets.

Within the address space, there are some address ranges which cannot be used for Virtual Networks, these are:

- 224.0.0.0 /4 (Multicast)
- 255.255.255.255 /32 (Broadcast)
- 127.0.0.0 /8 (Loopback)
- 169.254.0.0 /16 (Link-Local)
- 168.63.129.16 /32 (Internal DNS)

VNet

A virtual network is a representation of a network in the cloud. It enables resources in it to securely communicate with other resources within a virtual network. Multiple virtual networks can be created in a subscription. Each virtual network is isolated from other virtual networks.

The network architecture of ABC Company will follow a Hub and Spoke topology due to the presence of Shared components. The Shared Virtual network will act as a hub, and all other business unit virtual networks will be spokes. The Hub VNET will be peered with all other spoke VNETs. ABC Company currently does not have any cloud environment that makes use of hub and spoke topology.

Subnets

A subnet is an identifiably separate part of a virtual network. Azure VNETs supports the creation of 10,000 subnets.

Production subnets will be created resource wise (e.g., Jump box virtual machine, APIM) to ensure segregation that can be applied using NSG rules in subnets.

Non-Production subnets will also follow the same criteria as the production environment.

VNet Peering

Like a physical network, a Virtual Network (VNet) cannot communicate to other networks without a network router. As the term suggests, VNet Peering is used to connect two or more VNets in the same region to communicate with each other. The traffic routes through the internal Azure backbone in a VNet peering between subnets.

The following have been considered as important design characteristics:

- The Azure Infrastructure will be spanning across multiple virtual networks and subscriptions.
- Hub VNet hosted in Hub_Prod Subscription and Prod, Non-Prod-Dev, Non-Prod-Test VNET's hosted in Hub_NonProd Subscription respectively
- VNET peering will help link virtual network in the Hub and Spoke structure while providing guaranteed isolation aligned with the security standards.

IP Schema

The IP Schema will be based lined to cater for the customer's current requirements but can easily scale and be considered repeatable to allow further businesses to join the schema.

The Cloud hosted infrastructure will utilize the IP Address range reserved by ABC Company Networking team.

The IP Address Schema will correspond to customer standards, and it is assumed that it will be recorded by the Customer Networking team to avoid potential IP address conflicts when the Azure infrastructure is linked to On-Premises Data centers.

Network Topology

Network Topology refers to the layout of a network. How different nodes in a network are connected to each other and how they communicate is determined by the network's topology. It is the arrangement of different components in a network and plays a major role in every network architecture design. For ABC Company, the Hub & Spoke topology can be leveraged as shown below:

Figure 5 - Virtual Networks

Hub and Spoke Topology

Hub and Spoke topology are the system of connections arranged like a wire wheel in which all traffic moves along spokes connected to the hub at the center. Hub and spoke architectures are typically used for the environments which require segregation and central level management.

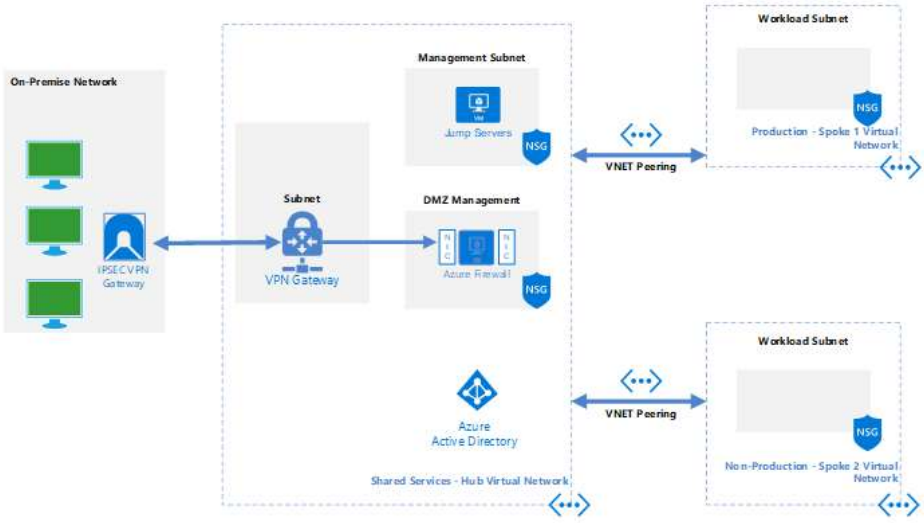

Figure 6 - Hub and Spoke Topology

The hub-spoke model provides the following benefits:

- Cost savings – a separate IPSEC VPN connectivity isn't required for each subscription. Shared/common services can also be hosted within the hub subscription, rather than being implemented in each spoke.
- Overcoming subscription limits – Microsoft Subscriptions have certain limits. By implementing a hub and spoke model, the business can have multiple subscriptions.

- Separation of concerns – Key IT services (i.e. Security and Identity) and can be centrally managed and controlled in the hub subscription which can provide environmental segregation.

Design Recommendation
ABC Company should utilize a hub-spoke design. The Hub will be Shared Service subscription which means any spoke services can utilize a common set of tools from the hub. An example of this would be an application server accessing authentication from the hub. **The Shared Services (Hub) will provide centralization and security for all workloads in the spoke subscriptions. This provides the following benefits:** • **Cost savings** • **Peering is kept to minimum** • **Environment level segregation** • **Future expansion & Common services adoption**

Azure Load Balancers

Azure Load Balancers can

- Scale-out applications to create high availability of services
- Support inbound and outbound scenarios that provide low latency and high throughput
- Scale-up to millions of flows for all TCP and UDP applications
- Distribute new inbound flows that arrive on frontend to backend pool instances, according to rules and health probes
- Provide outbound connections for virtual machines (VMs) inside a virtual network by translating their private IP addresses to public IP addresses.

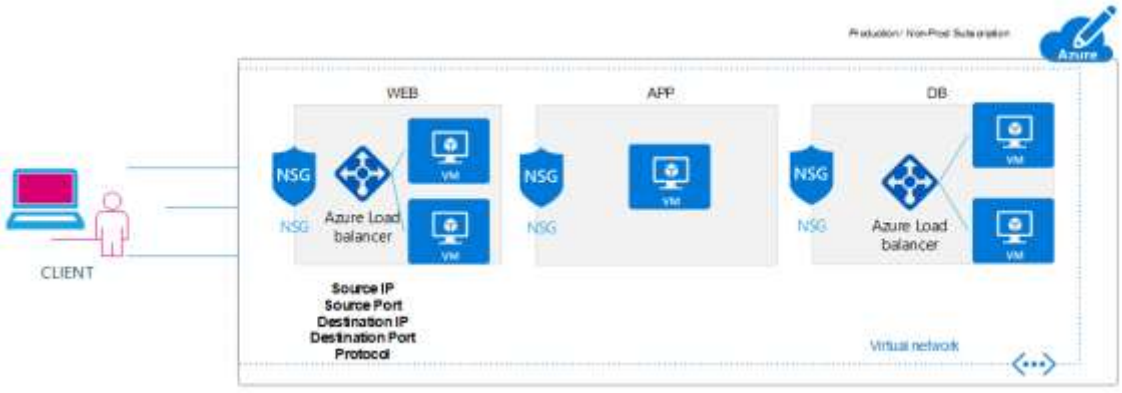

Figure 7 - Load Balancer

An internal Load Balancer as shown in **Figure** above in the business tier directs traffic only to resources that are inside a virtual network or that use a VPN to access Azure infrastructure. In this respect, an internal Load Balancer differs from a public Load Balancer.

Azure infrastructure restricts access to the load-balanced frontend IP addresses of a virtual network. Frontend IP addresses and virtual networks are never directly exposed to an internet endpoint. Internal line-of-business applications run in Azure and are accessed from within Azure or from on-premises resources.

An internal Load Balancer enables the following types of load balancing:

- Within a virtual network: Load balancing from VMs in the virtual network to a set of VMs that reside within the same virtual network.
- For a cross-premises virtual network: Load balancing from on-premises computers to a set of VMs that reside within the same virtual network.
- For multi-tier applications: Load balancing for internet-facing multi-tier applications where the backend tiers are not internet-facing. The backend tiers require traffic load-balancing from the internet-facing tier.
- For line-of-business applications: Load balancing for line-of-business applications that are hosted in Azure without additional load balancer hardware or software.

Azure Load Balancer is available in two SKUs:

- Basic Load Balancer
- Standard Load Balancer

Feature	Standard Load Balancer	Basic Load balancer
Backend pool size	Supports up to 1000 instances.	Supports up to 100 instances.
Health probes	TCP, HTTP, HTTPS	TCP, HTTP
Backend pool end points	Any virtual machine in a single virtual network, including blend of virtual machines, availability sets, virtual machine scale sets.	Virtual machines in a single availability set or virtual machine scale set.
Availability Zones	In Standard SKU, zone-redundant and zonal frontends for inbound and outbound, outbound flows mappings survive zone failure, cross-zone load balancing.	Not available.
HA Ports	Internal Load Balancer	Not available.

SLA	99.99% for data path with two healthy virtual machines.	NA
Pricing	Charged based on number of rules, data processed inbound and outbound associated with resource.	No charge

Table 5 - Load Balancing Options

Note

It is impossible to move the IP addresses associated with existing Basic Load Balancer seamlessly to Standard Load Balancer since they have different SKUs. Re-configuration with downtime would be required.

Design Recommendation

As there is sensitive data for ABC Company applications, we should configure the environment with "Standard Load Balancers" as part of each subnet.

We should implement Load balancers for each subnet (Web, App and DB subnets) for the Production spoke at a minimum.

Azure Connectivity

For this design there are two forms of IPSEC VPN that should be established for connectivity in and out of Azure. Both Services will terminate via an Azure Traffic Manager that will load balance the traffic to two VPN Gateway's to manage the traffic flows in the environment.

Site to Site VPN

A Site-to-Site VPN connection is used to connect on-premises network to an Azure virtual network over an IPsec/IKE (IKEv1 or IKEv2) VPN tunnel. This type of connection requires a VPN device located on-premises that has an externally facing public IP address. Site to Site VPN connection should be established from ABC Company data centers to Microsoft Azure.

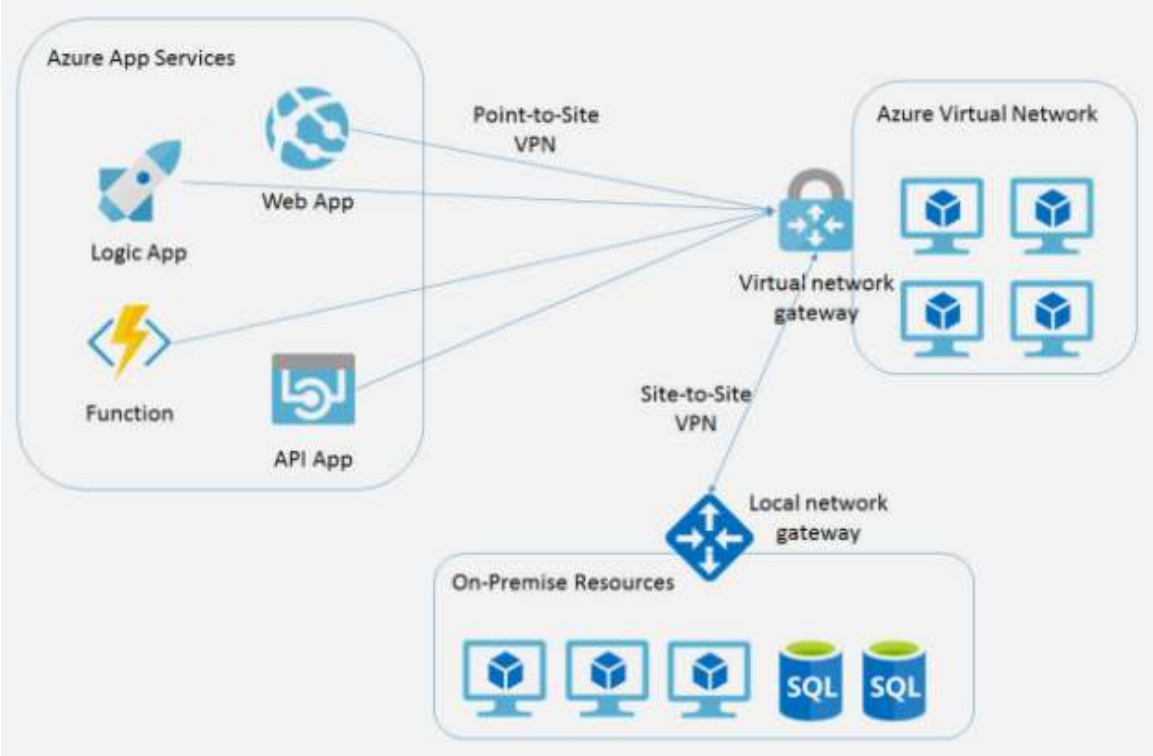

Figure 8 - VPN Connectivity

Azure Traffic Manager

Azure Traffic Manager is used to manage traffic from multiple incoming connections at one given time. The Azure Traffic Manager will pass requests to the least busy Azure VPN Gateway to handle IPSEC VPN connections. Azure Traffic Manager operates at the DNS layer to quickly and efficiently direct incoming DNS requests based on the routing method. An example would be sending requests to the closest endpoints to improve the responsiveness of your applications.

Design Recommendation
The Azure Traffic Manager should be configured to monitor the workloads and respond based on utilization, weight, and geography. The Azure Traffic Manager should also pass "Route-based" connectivity parameters for on-premises connectivity. Depending on business requirements, ABC Company can choose a higher or lower probing frequency to switch between on-premises to Azure in a disaster event, and ensure minimal downtime for users as shown in the example below.

On-premises to Azure failover

In a disaster event, a company can trigger a failover to Azure and recover its applications on Azure. The Priority traffic-routing method in Azure Traffic Manager allows a company to easily implement a failover pattern such as the below.

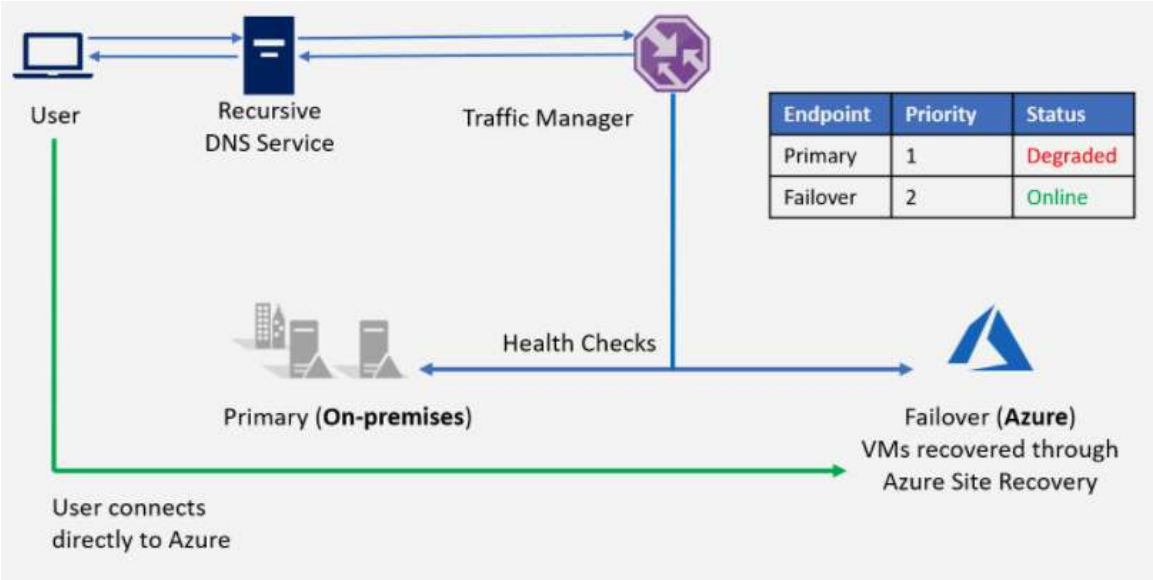

Figure 9 - Failover

Security

Network Security Group

Network Security Groups ("NSG") are a key part of controlling network access between resources in Azure. They act as mini firewalls that define rules of what traffic can get from one resource to another and even to/from the internet. Network Layer security using ("NSGs") should be enabled at subnet level. NSG rules contain the following properties:

Property	Description
Name	Name of the Rule.
Protocol	Protocol to match with the rule.
Source port range	Source port range to match with the rule.
Destination port range	Destination port range to match with the rule.
Source address prefix	Source address prefix or tag to match with the rule.
Destination address prefix	Destination address prefix or tag to match with the rule.
Direction	Direction of traffic to match the rule.
Priority	Rules are checked in the order of priority.

Table 6 - NSGs

Design Recommendations – Inbound and Outbound Default Rules

Default NSG rules should be applied to all subnet levels with:

- **All outbound traffic to internet will be blocked**
- **NSG rules should be defined within the same** VNet **and between** VNet**s**
- **NSG rules should be defined for enabling internet connectivity on a need's basis**

Inbound NSG Rules

Name	Priority	Source IP	Source Port	Destination IP	Destination Port	Protocol	Access
Allow VNET Inbound	65000	Virtual Network	*	Virtual Network	*	*	Allow
Allow Azure Load Balancer Inbound	65001	Azure Load Balancer	*	Virtual Network	*	*	Allow
Deny All inbound	65500	*	*	*	*	*	Deny

Outbound NSG Rules

Name	Priority	Source IP	Source Port	Destination IP	Destination Port	Protocol	Access
Deny Internet Outbound	4000	*	*	Internet	*	*	Deny
Allow VNet outbound	65000	Virtual Network	*	Virtual Network	*	*	Allow
Allow Internet Outbound	65501	*	*	Internet	*	*	Allow
Deny all outbound	65500	*	*	*	*	*	Deny

Azure Firewall

Azure Firewall is a managed, stateful, cloud-based network security service that protects Azure Virtual Network resources.

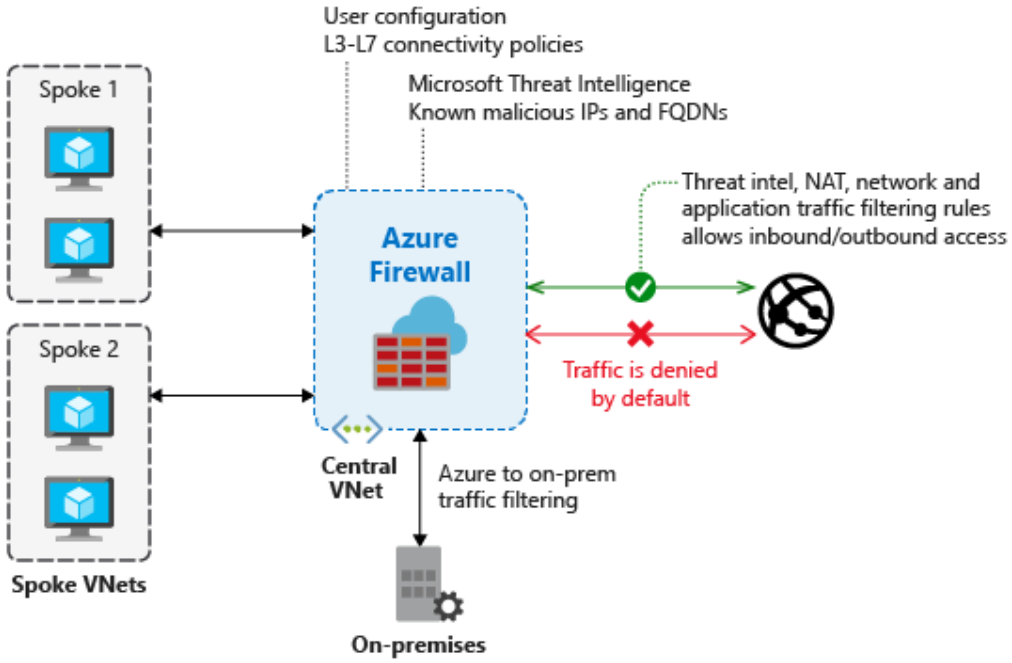

Figure 10 - Azure Firewall

Azure Firewall uses a static public IP address for the virtual network resources allowing outside firewalls to identify traffic originating from the virtual network. The service is fully integrated with Azure Monitor for logging and analytics.

Azure Firewall offers the following features:

- Built-in high availability
- Availability Zones
- Unrestricted cloud scalability
- Network traffic filtering rules
- Threat intelligence

Design Recommendation
Azure Firewalls should be utilized • One per region to handle requests from each Azure VPN Gateway • Should reside in the Hub Shared Services subscription in the DMZ for internal and external connections

DDoS Protection

Distributed denial of service (DDoS) attacks are a serious availability and security concern facing customers that are moving their applications to the cloud. A DDoS attack attempts to exhaust an application's resources, making the application unavailable to legitimate users by making continuous fake requests to the server or application. DDoS attacks can be targeted at any endpoint that is publicly reachable through the internet. Azure DDoS Protection along with application best practices can provide protection from DDoS attacks.

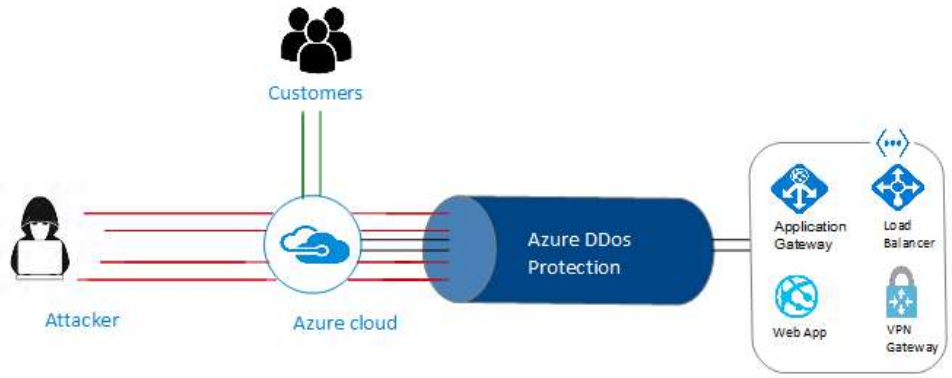

Figure 11 - DDOS

Azure DDoS protection has basic and standard service tiers:

Basic - Automatically enabled as part of the Azure platform. Always-on traffic monitoring, and real-time mitigation of common network-level attacks, provide the same defenses utilized by Microsoft's online services.

Standard - Provides additional mitigation capabilities over the basic service tier that are tuned specifically to Azure Virtual Network resources. DDoS Protection Standard is simple to enable and requires no application changes. Protection policies are tuned through dedicated traffic monitoring and machine learning algorithms. Policies are applied to public IP addresses associated with resources deployed in virtual networks, such as Azure Load Balancer, Azure Application Gateway, and Azure Service Fabric instances, but this protection does not apply to App Service Environments. Real-time telemetry is available through Azure Monitor views during an attack, and for history.

DDoS Protection Standard can mitigate the following types of attacks:

- Volumetric attacks: Volumetric attacks are the most common type of DDoS attack. Volumetric attacks are brute-force assaults that target the network and transport layers. The attack's goal is to flood the network layer with a substantial

amount of seemingly legitimate traffic. It includes UDP floods, amplification floods, and other spoofed-packet floods.

- Protocol attacks: These attacks render a target inaccessible, by exploiting a weakness in the layer 3 and layer 4 protocol stack. It includes SYN flood attacks, reflection attacks, and other protocol attacks.
- Resource (application) layer attacks: These attacks target web application packets, to disrupt the transmission of data between hosts. The attacks include HTTP protocol violations, SQL injections, cross-site scripting, and other layer 7 attacks.

DDoS Protection Standard features include:

- Native platform integration
- Turn-key protection
- Always-on traffic monitoring
- Attack analytics
- Attack metrics
- Attack alerting

DDoS Protection Standard monitors actual traffic utilization and constantly compares it against the thresholds defined in the DDoS Policy. When the traffic threshold is exceeded, DDoS mitigation is initiated automatically. When traffic returns below the threshold, the mitigation is removed. During mitigation, traffic sent to the protected resource is redirected by the DDoS protection service and several checks are performed.

DDoS protection blocks attack traffic and forwards the remaining traffic to its intended destination. Within a few minutes of attack detection, notification is being sent using Azure Monitor metrics. By configuring logging on DDoS Protection Standard telemetry, logs can be written to be available giving options for future analysis. Metric data in Azure Monitor for DDoS Protection Standard is retained for 30 days.

Design Recommendation
DDoS Protection “resource” should be created to protect from distributed denial of service (DDoS) attacks for each subscription with always-on monitoring and automatic network attack mitigation. The basic plan should be used. There will be no up-front commitment, and cost will scale with customer cloud deployment. Azure DDoS service will be integrated with all Virtual Networks (VNETs) and will provide protection for Azure applications from any impacts relating to DDoS attacks.

Azure Information Protection (Data Security)

In Azure, there are several approaches to be considered when addressing data protection. The following section describes how to use these approaches:

- Data Classification
- Protecting data in transit
- Protecting data at rest

Data Classification

Data classification provides a way to categorize organizational data based on different levels of criticality. The data classification process categorizes data by sensitivity and business impact to identify risks. Once data is classified, it can be managed in ways that protect sensitive or important data from theft or loss.

Design Recommendation

The following classification should be implemented as part of the Greenfield deployment:

- Non-business: Data from personal life that does not belong to enterprise.
- Public: Business data that is freely available and approved for public consumption.
- General: Business data that is not meant for a public audience.
- Confidential: Business data that could cause harm to enterprise if overshared.
- Highly confidential: Business data that would cause extensive harm to ABC Company if overshared.

Note

These classifications can be changed should the business requirements change.

It is recommended that the Azure resources are tagged that will hold such data. But tagging cloud assets by classification is not a replacement for a formal data classification process, it only provides a valuable tool for managing resources and applying policy.

Protecting data in transit

Encryption of data in transit is a mechanism of protecting data when it is transmitted across networks in Azure environment. Within Azure Storage, data in transit can be secured using transport level encryption (Such as HTTPS) when data is transferred in/out of Azure storage. HTTPS protocol ensures secure communication over the public Internet.

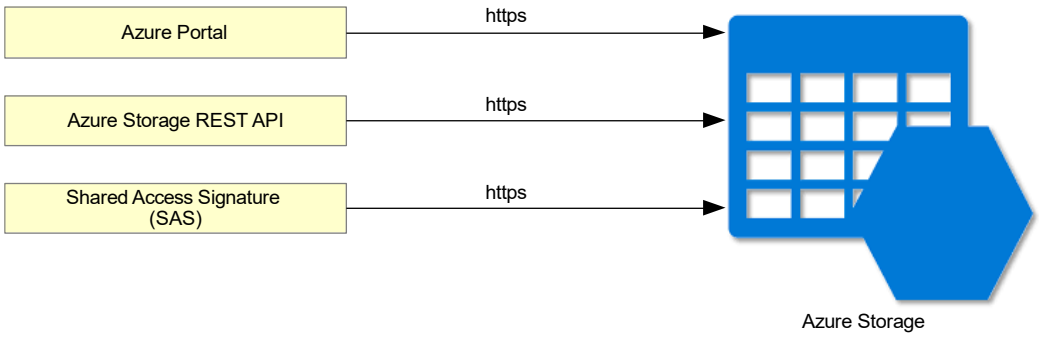

Figure 12 - Data Encryption

Storage Account – Data in Transit

Shared Access Signatures can be used to delegate access to Azure Storage objects. A shared access signature is a signed Uniform Resource Identifier ("URI") that points to one or more storage resources and includes a token that contains a special set of query parameters. There is an option to specify that only the HTTPS protocol can be used when using Shared Access Signatures, ensuring that anybody sending out links with SAS tokens will use the proper protocol. The token indicates how the resources may be accessed by the client (i.e. read, write, delete, etc.)

Design Recommendation
An Azure Storage Account should be provided as part of the shared subscription model to each of the spoke subscriptions. The Azure storage account should have "Shared Access Signatures" generated from the Blob containers under the security parameters configured using "Azure Storage Explorer". Each Shared Access signature shall be unique to protect the data in transit and will have "access policies" configured with permission to allow read, write, delete or list and will utilize Azure AD for storage endpoint domain authentication. This will create and URL with SAS to be used as appropriately.

Protecting data at rest

Encryption of data at rest is the encoding (encryption) of data when it is persisted. Encryption at rest is designed to prevent the attacker from accessing the unencrypted data by ensuring the data is encrypted when on disk. If an attacker obtains a hard drive with encrypted data but not the encryption keys, the attacker must defeat the encryption to read the data.

Data at rest can be encrypted using below methods:

- **Storage Service Encryption (SSE):** This encryption technique is an inbuilt feature of Azure; data will be encrypted prior to its entry into Storage Account. Decryption will take place prior to its retrieval from Storage Account. Storage Service Encryption uses 256-bit Advanced Encryption Standard (AES) encryption.

Encryption is enabled at Storage Account level, hence irrespective of data classification all data in the Storage Account will be encrypted.

- **Azure Disk Encryption:** Azure Disk Encryption leverages the industry standard BitLocker feature of Windows and the DM-Crypt feature of Linux to provide volume encryption for the OS and data disks in Azure VM. This solution is integrated with Azure Key Vault, hence can control and manage the disk-encryption keys used in Azure VM's and can audit their usage in Azure key vault. Azure Disk Encryption also ensures that all data on the Azure VM disks are encrypted at rest in Azure storage. Encryption and Decryption of data adds an additional latency overhead for read write operations.

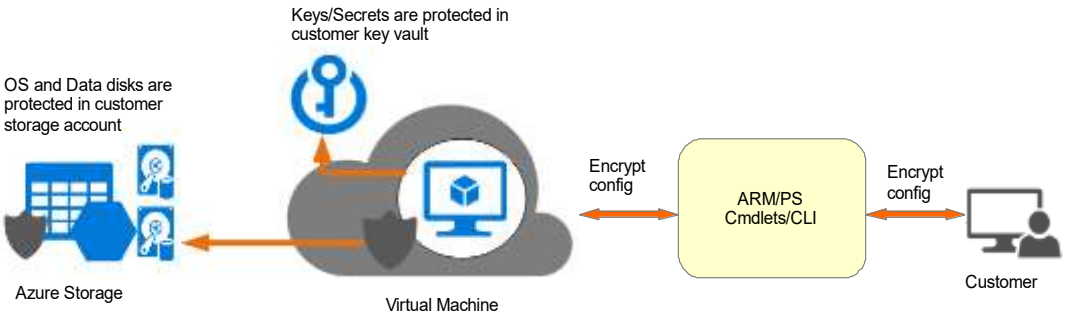

Figure 13 - Azure Key Vaults

Design Recommendation
The native Azure Service of "Storage Service Encryption" for all spoke storage services should be utilized.

Azure Advisor

Azure Advisor is a personalized cloud consultant that helps to follow best practices to optimize the Azure deployments. It analyses resource configuration and usage telemetry and then recommends solutions that can help to improve cost effectiveness, performance, high availability, and security of the Azure resources.

The Advisor dashboard can display personalized recommendations for all the subscriptions. The customer team can apply filters to display recommendations for specific subscriptions and resource types in the azure portal.

The recommendations of azure advisor are divided into four categories:

1. High Availability
2. Security
3. Performance

4. Cost

Design Recommendation
The recommendation for the above have been stipulated by Microsoft at: https://docs.microsoft.com/en-us/azure/advisor/advisor-overview The Azure Advisor will report out of the box on the above and shall help the customer understand the categories that are configured. When detailed requirements are provided, we can assist with further personalized recommendations for each of the subscriptions.

Azure Sentinel (SIEM)

Microsoft Azure Sentinel is a scalable, cloud-native, security information event management (SIEM) and security orchestration automated response (SOAR) solution. Azure Sentinel delivers intelligent security analytics and threat intelligence across the enterprise, providing a single solution for alert detection, threat visibility, proactive hunting, and threat response.

Working across on-premises and in-cloud infrastructure, it's intended to be easy to set up, low maintenance, and easy to use. By building on cloud-scale data collection, and on Microsoft's own threat detection tools, Azure Sentinel can automate response using orchestration across your entire estate. It's software-as-a-service so it's scalable, and you only pay for the resources you use. The data for this analysis is stored in an Azure Monitor Log Analytics workspace. Azure Sentinel is billed based on the volume of data ingested for analysis in Azure Sentinel and stored in the Azure Monitor Log Analytics workspace. Features include

- Collect data at cloud scale.
- Detect previously undetected threats.
- Investigate threats with artificial intelligence.
- Respond to incidents rapidly.

Azure Sentinel is a hybrid cloud security solution, capable of processing and analyzing data from Azure and other cloud provider services, as well as from Windows and Linux workloads no matter whether they are on-premises or in a cloud. There are two ways to pay for the Azure Sentinel service:

Capacity Reservations

With Capacity Reservations you are billed a fixed fee based on the selected tier, enabling a predictable total cost for Azure Sentinel. Capacity Reservation provides you with a discount (up to 60%) on the cost based on your selected capacity reservation compared to the Pay-As-You-Go pricing. You have the flexibility to opt out of the capacity tier any time after the first 31 days of commitment.

Pay-As-You-Go

With Pay-As-You-Go pricing, you are billed per gigabyte (GB) for the volume of data ingested for analysis in Azure Sentinel and stored in the Azure Monitor Log Analytics workspace.

Design Recommendation
We recommend using Azure Sentinel to be consumed as a "Capacity Reservation" as an Enterprise Agreement is secured with this deployment. The Azure Sentinel should be consumed from the Azure marketplace and configured with the Shared Services Subscription. Should the customer not require this service after 31 days, then the service can be ceased. The minimum commitment is 31 days from activation.

Azure Security Center (End Point Protection)

Azure Security Center is a unified infrastructure security management system that provides advanced threat protection across workloads in the Azure. PaaS services in Azure including Service Fabric, SQL databases, and storage accounts are monitored and protected by Security Center without necessitating any deployment.

The features and benefits of Azure Security Center is detailed below:

Feature	Benefit
Strengthening security posture	• ABC Company can set their policies to run on management groups, across azure subscriptions, and even for a whole azure tenant. • Can optimize and improve security by configuring recommended controls in ABC Company azure resources. • Security Center will continuously discover new resources that are being deployed across ABC Company azure subscription and assesses whether they are configured according to security best practices.
Protect against threats	• Will enable ABC Company to detect and prevent threats at the Infrastructure as a Service (IaaS) layer, as well as for Platforms as a Service (PaaS) in ABC Company Azure environment.
Get Secure Faster	• Native Azure integration (including Azure Policy and Azure Monitor logs) combined with other Microsoft security solutions, such as Microsoft Cloud App Security and Windows Defender Advanced Threat Protection help make sure ABC Company security solution is comprehensive as well as simple to onboard and roll out. • ABC Company can pull together complete security policies including Azure Policy and built-in Security Center policies across all ABC Company Azure resources.

Table 7 - Security Center

Design Recommendation
Azure Security Center should be provided from the Shared Service Subscription model to detect and prevent threats across all spoke services. ABC Company should implement the out-of-the-box configuration to allow basic monitoring within the Security center. By default, the following should be configured: • Virtual Machine Behavioral Analytics threat detection alerts • File less threat detection alerts • Network-based threat detection alerts • File Integrity Monitoring • Network map • Regulatory Compliance dashboard & reports • Endpoint protection assessment • Network security assessment

Azure Virtual Machines

Virtual machines (VMs)

Azure Virtual machines are the resource, which gives the flexibility of virtualization for a wide range of computing solutions with support for windows server, Linux, SQL server, Oracle more. All the current generation Virtual machines include load balancing and auto scaling. Virtual machines with managed disks will provide optimal performance. It is the Microsoft recommended feature.

Azure virtual machines are in Pay-As-You-Go model. We need to pay only for the compute capacity with no long-term commitment or upfront payments. We can increase or decrease the compute capacity on demand. We can start or stop the virtual machines at any time and pay only for the consumption.

The disadvantage associated with the Pay-As-You-Go model is deep discounts. At ABC Company, the interdependencies between the Production, UAT and Development workloads do not provide the ability to shut down enough images (Part-time) to take advantage from lowered utilization in this consumption model.

From the data collection and assessments data, the 1-year Azure reserved instances model provides ~41% overall savings. While 3-year commitments can save ~53%, rapid technology change and competition have led to better pricing sooner.

Design Recommendation
We recommend using Azure Subscriptions that utilizes Azure Reserved Instances for a 1-year minimum commitment for all workloads at ABC Company.

Cost of Compute in a Region

A common point of contention is that each region has different pricing. The same server in one region can be priced differently in another region. In fact, on azureprice.net you can see why below we chose East US 2/Central US.

Figure 14 - Regional Pricing

VM Management and Monitoring

From a management and monitoring perspective, the Azure virtual machines can utilize Operations Management Suite ("OMS"). Log Analytics is also known for OMS, which provides insights. These insights include information on virtual machines including, but not limited to:

- CPU utilization
- Memory utilization
- Network utilization
- Windows update
- Event Viewer monitoring

These VM metrics are displayed in Log Analytics through dashboard and title views. VMs that are created outside of the Self-Service Portal UI must have the Log Analytics agent extension installed and be registered in the Log Analytics workspace. VMs should only register to the Log Analytics workspace that is associated with their subscriptions.

Design Decision
ABC Company should implement the above utilization metrics for all workloads in the associated spokes.

High Availability - Availability Set

For applications and services that have high availability requirements, we will deploy multiple VMs hosting the application in an availability set. For High Availability, we recommend the following guidance:

- Ensure that the application can handle single-server outages.
- For best results ensure each application tier is in its own availability set.
- VMs in an availability set must be in the same resource group.
- An availability set can have a maximum of 100 VMs.

Design Decision
The customer should implement the “High Availability” feature, if there is requirement from the business.

Storage

Virtual Machine Storage

We recommend the following general guidance for VM storage

- Use managed disks for all VMs
    - Managed Disks simplify disk management for Azure IaaS VMs by handling storage account management.
    - Only Locally Redundant Storage (LRS) is available for Managed Disks.
- For situations which require unmanaged disks:
    - Use Azure Availability Sets when deploying systems that will support redundant services.
    - Use Locally Redundant Storage (LRS) storage accounts for VMs (LRS storage accounts have higher ingress and egress bandwidth than the other types) unless using Premium Storage.
    - Create storage accounts specifically for VHDs and keep no more than approximately 40 active VHDs in each VHD account.

    - Use multiple storage accounts for VMs in an Availability Set to avoid the scenario of a single storage account going offline impacting all systems in an Availability Set.

- Use premium storage for all production workloads (especially databases) and critical infrastructure.
    - Azure premium SSDs deliver high-performance and low-latency disk support for virtual machines (VMs) with input/output (IO)-intensive workloads.
    - Premium SSD can provide twice the IOPS of standard SSD.
    - Premium storage guarantees the capacity and IOPS of the disk.
    - Premium SSD disks are designed to provide low single-digit millisecond latencies and target IOPS and throughput.
- To achieve VM single-disk capacity greater than 1TB create multiple VHDs and create a stripe set in the OS. However, beware of the limitations of this approach particularly on backup and replication strategies. GRS and any snapshot/backup are not guaranteed to be volume consistent.
- Do not use Storage Service Encryption (SSE) to encrypt the VM/VHD-specific storage accounts. Use whole-disk encryption solutions (Bit locker/DM-crypt) where appropriate instead.

Property	Premium SSD	Standard SSD	Standard HDD
Recommended Scenarios	Production and performance sensitive workloads	Web servers, lightly used enterprise applications and dev/test	Backup, non-critical, infrequent access
Max disk size	32,767 Gibb	32,767 GiB	32,767 GiB
Max throughput	900 MiB/s	750 MiB/s	500 MiB/s
Max IOPS	**20,000**	**6000**	**2000**

Table 8 - Disk Storage Options

Management of storage accounts

- The IOPS limitation of storage accounts presents a management challenge. A Standard storage account has a maximum total request rate of 20,000 IOPS. Storage accounts must be monitored regularly to ensure that the account is not over the IOPS limit.

- Azure Monitor can monitor IOPS of Azure storage account by PercentThrottlingError metrics. Throttling errors occur when a storage service exceeds its target scalability. If the PercentThrottlingError metric shows an increase in the percentage of requests that are failing with a throttling error, it can be two different scenarios.
    - Transient increase in PercentThrottlingError
    - Permanent increase in Percent Throttling Error
- Azure can set alerts for these metrics to notify the customer whenever throttling events occur.

Design Decision
This option is activated by default. ABC Company should monitor the alerts in Azure Monitor.

Azure Storage Encryption

Azure Storage automatically encrypts the data when persisting it in the cloud. SSE provides encryption-at-rest and safeguards the data to meet the organizational security and compliance commitments. It is enabled by default for all Managed Disks, Snapshots, and Images.

Azure Storage Service Encryption (SSE) is compatible with ASR and provides protection and safeguards data to meet organizational security and compliance commitments. With this feature, Azure Storage automatically encrypts the data prior to persisting in storage and decrypts prior to retrieval without any performance impact or visibility to end users.

Design Decision
This option is activated by default. ABC Company should monitor the alerts in Azure Monitor.

Disaster Recovery

Backups

Azure Backup is a Built-in PaaS Service that can be used to protect and restore Azure virtual machines (VMs), and On-premises machines and workloads. It uses Recovery Services Vault for storing data in the cloud. A vault is an online-storage entity used to

hold data such as backup copies, recovery points, and backup policies. When the backup job for a protected resource run, it creates a recovery point inside the Recovery Services vault. One of these recovery points can be used to restore data to a given point in time. There is provision to create and use policies to define when a backup job runs and how long the recovery points are stored.

Azure Backup offers two types of replications to keep storage/data highly available.

- Locally redundant storage (LRS): Replicates data three times in a storage scale unit in a data center (same region).
- Geo-redundant storage (GRS): Default and recommended replication option. GRS replicates data to a secondary region (hundreds of miles away from the primary location of the source data).

Design Decision
ABC Company should implement Geo-Replication Storage Replication with 30 days of retention.

Disaster Recovery Standards

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. Across the regional pairs, Azure serializes platform updates (planned maintenance), so that only one paired region is updated at a time. In the event of an outage affecting multiple regions, at least one region in each pair will be prioritized for recovery.

Geography	Paired Regions	
US	East US 2 (Primary)	Central US (Secondary)
Canada	Canada East (Primary)	Canada Central (Secondary)

Table 9 - Azure Paired Regions for DR

Primary Region: Primary Region is the primary site of workload deployment for ABC Company Azure Cloud Platform

Secondary Region: Secondary Region also known as the DR Region, where the workloads will be hosted based on ABC Company DR strategy in case of primary Region failure.

Backup Considerations

Applications should be ranked based on business severity and DR plan should be based on the application severity.

Tier	Business Severity	RTO	RPO	DR Strategy	Cost Index
1	Very High	30 minutes	<1 minute	Active-Active	Very High
2	High	30 minutes	15 minutes	Active-Active	High
3	Moderate	2 hours	30 minutes	Warm Standby	Moderate
4	Low	1 day	12 hours	Pilot Light	Low
5	Non-Critical	2 days	1 day	Cold Site (Backup & Restore)	Very Low

Table 10 - Azure Backup Options

The following file types will be backed up based on the backup policy for recovery purposes on both primary and secondary Azure Regions.

- VM backups
- VM Disk snapshots
- Important configurations
- Critical system & service logs
- Critical Database backups & logs

The frequency of backup operations should be based on the RPO value defined for the application. The different types of files that are backed up in the primary Azure Region will periodically be copied to the secondary Azure Region in an asynchronous manner or will need to be made available at secondary region using Geo Redundancy and Azure Site Recovery.

Design Decision
ABC Company should detail backup requirements as above.

Azure DNS

DNS domains in Azure DNS are hosted on Azure's global network of DNS name servers. Each DNS query is answered by the closest available DNS server to provide fast performance and high availability for domains. Azure DNS can leverage Azure AD,

auditing, governance, role-based access control (RBAC), and resource locking to secure DNS service.

Design Decision
ABC Company should configure DNS services within Azure but should also provide external DNS forwarders to allow traffic to be routed to Azure.

Azure Site Recovery

ASR helps ensure business continuity by keeping business apps and workloads running during outages. Site Recovery replicates workloads running on physical and virtual machines (VMs) from a primary site to a secondary location. When an outage occurs at primary site, you fail over to the secondary location to access apps and resources. After the primary location is running again, you can fail back.

Azure Site Recovery can be used for:

- Disaster recovery of Azure VMs from a primary region to a secondary region
- Replicate on-premises VMs and physical servers to Azure, or to a secondary on-premises data center.

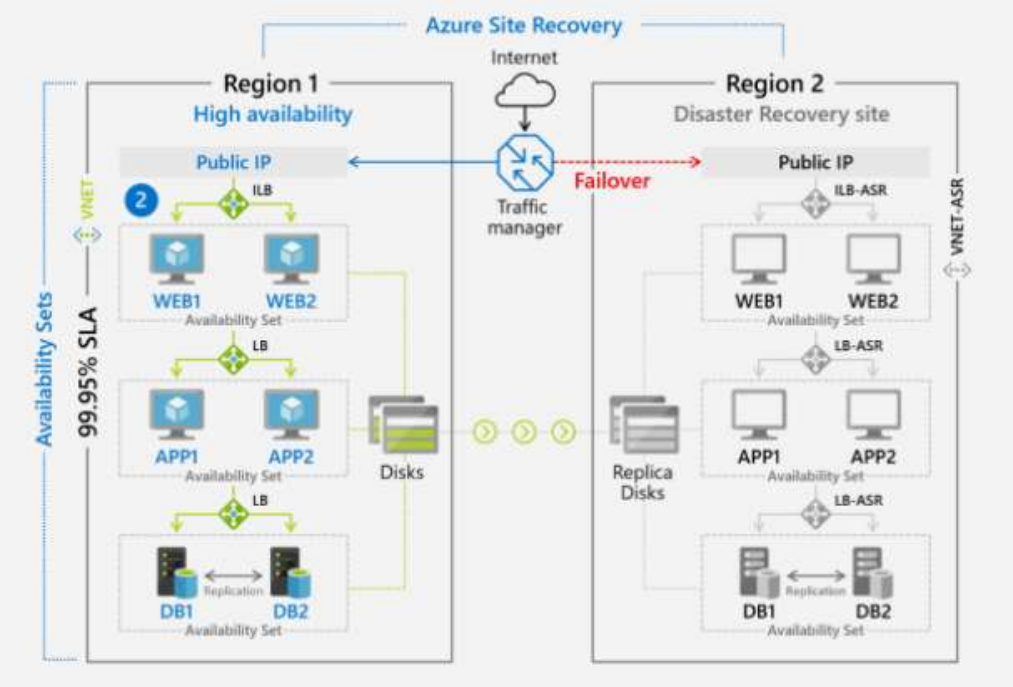

Figure 15 - Site Recovery

To enable replication of a VM and essentially copy data by using Azure Site Recovery, user must have:

- Permissions to create a VM in Azure resources. The Virtual Machine Contributor built-in role has these permissions, which include:
  - Permission to create a VM in the selected resource group
  - Permission to create a VM in the selected virtual network
- Permission to write to the selected storage account

Permissions to manage Azure Site Recovery operations. The Site Recovery Contributor role has all the permissions that are required to manage Site Recovery operations in a Recovery Services vault.

Design Decision
As mentioned earlier, ABC Company should provide detailed backup requirements to build out this architecture.

Identity Management

Azure Active Directory & Domain Services

Azure Active Directory will be used for authenticating User Logins and Azure Active Directory Domain Services will be used to join the VMs to Active Directory Domain.

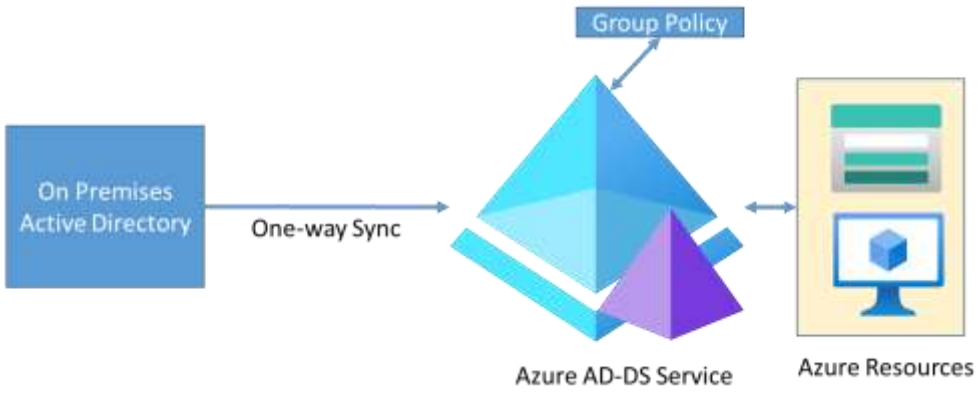

Figure 16 - On-premises and AD-DS Connectivity

Design Decision
VMs will be joined to the AADDS Domain for Centralized management and Authentications. Azure AD users should be synced to AADDS.

User Management

Azure AD administrators can perform identity management tasks for users in terms of groups, licenses, deployed enterprise apps, and administrator roles. As the organization grows, you can use Azure AD groups and administrator roles to:

- Assign licenses to groups instead of to individually.
- Delegate permissions to distribute the work of Azure AD management to less privileged roles
- Assign enterprise app access to groups.

Role name	Permissions summary
Application Administrator	Can add and manage enterprise applications and application registrations and configure proxy application settings. Application Administrators can view Conditional Access policies and devices but not manage them.
Cloud Application Administrator	Can add and manage enterprise applications and enterprise app registrations. This role has all the permissions of the Application Administrator, except it can't manage application proxy settings.
Application Developer	Can add and update application registrations but can't manage enterprise applications or configure an application proxy.

Design Decision
ABC Company should use the built-in resource groups to provide permissions across the services.

Role Based Access Control (RBAC)

Azure Role-Based Access Control (RBAC) offers access management of Azure resources so that users will get adequate permission to perform work tasks for their assigned role(s). Azure RBAC:

- Allows secure access with granular permissions
- Is it assignable to users, groups, or service principals
- Has built-in roles that make it easy to get started
- Allows the additional flexibility of custom roles
- Also has three basic roles that apply to all resource types:
    - Owner has full access to all resources including access delegation to others.
    - Contributors can create and manage all types of Azure resources but cannot grant access to others.

- Readers can view existing Azure resources

Many other built-in RBAC roles are available to manage or administer specific Azure resources. In addition, Azure allows creation of custom RBAC roles to meet specific access requirements. Custom roles can be assigned to users, groups, and service principals at subscription, resource group, and resource scopes. Custom roles are stored in an Azure Active Directory (Azure AD) directory and can be shared across subscriptions. Each directory can have up to 5000 custom roles.

Role Based Access Control (RBAC) is configurable at different levels in Azure. The most restrictive role is the Global-Administrator role and should be completely restricted.

Design Decision
ABC Company should use the built-in resource groups to provide permissions across the services.

Management & Operations

Infrastructure Monitoring

Monitoring services are used to analyze the performance and activity logging of various Azure resources. Azure Monitor is Microsoft's built-in monitoring service. With a wide array of metrics, Azure monitor visualizes data about the throughput of Azure services such as VM, storage, databases, etc.

Azure Monitor provides basic alerting capabilities that can generate alerts based on the thresholds configured for a given metric.

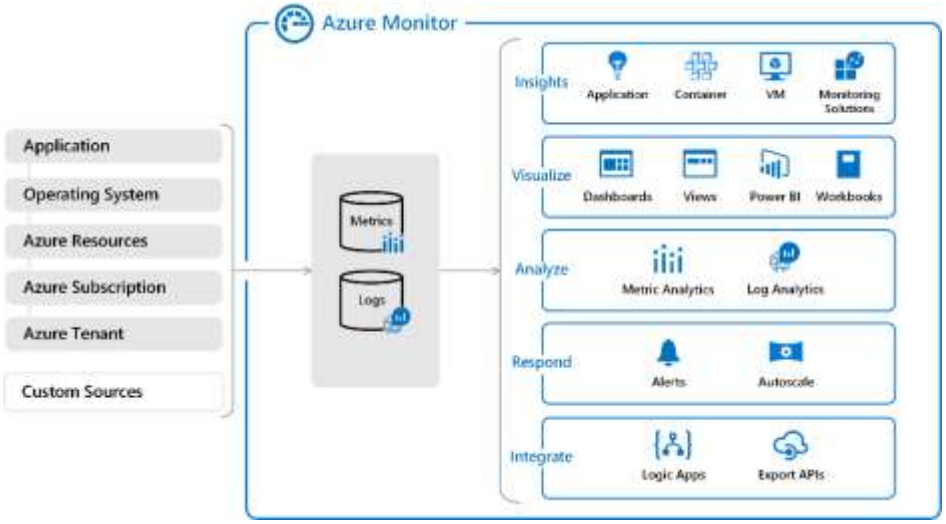

Figure 17 - Azure Monitoring

Network Monitoring

Traffic Analytics is an Azure native solution that provides visibility to users and application activity in cloud networks. Azure virtual networks have NSG flow logs, which provide information about ingress and egress IP traffic through a Network Security Group associated to individual network interfaces, VMs, or subnets.

Traffic Analytics can visualize network activity across the Azure subscriptions and identify hot spots. It can also identify security threats and provide information such as open-ports, applications attempting internet access, and virtual machines (VM) connecting to rogue networks. Some of the prerequisites for using Traffic Analytics are as follows:

- An Azure Log Analytics workspace, with read and write access.
- An Azure Storage account, to store raw flow logs.
- A Network Watcher enabled subscription.
- Network Security Group (NSG) flow logs enabled for the NSGs that need to be monitored.

Decision
ABC Company should configure Network Monitoring in the Shared Services Hubs.

Automation

Azure Automation delivers a cloud-based automation and configuration service that provides consistent management across Azure environments.

Azure Automation Account - Azure Automation account/s can be created through Azure. This method provides a browser-based user interface for creating and configuring Automation accounts and related resources.

Update Management - Update Management solution can be used in Azure Automation to manage operating system updates for Windows and Linux computers in Azure and on-premises environments. We can quickly assess the status of available updates on all agent computers and manage the process of installing required updates for servers. Update Management can be enabled for virtual machines (VMs) directly from your Azure Automation account.

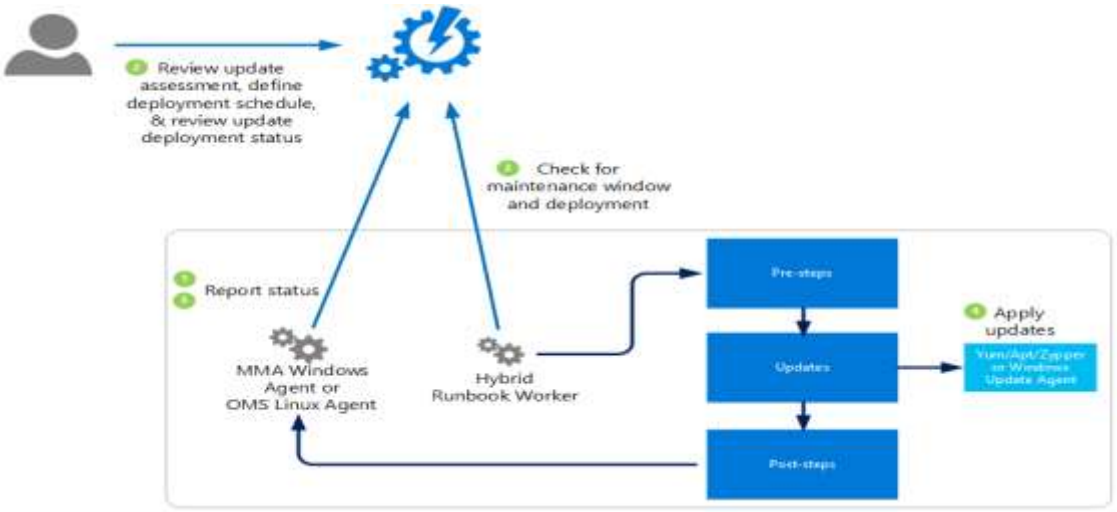

Figure 18 - : Update Management

Design Decision
The update management service should be implemented to perform updates directly in availability sets for HA/DR.

Microsoft Azure DevOps

Development and Testing Automation also comes in the form of Microsoft Azure DevOps, a hosted service providing development and collaboration tool that was formerly known as Visual Studio Team Services (VSTS).

In 2018, Microsoft split VSTS into 5 different Azure Branded Services, under the banner Azure DevOps for a comprehensive offering in the Public Cloud that makes it easier for developers to adopt portions of the Azure DevOps platform, without requiring them to go "all in" like the former VSTS.

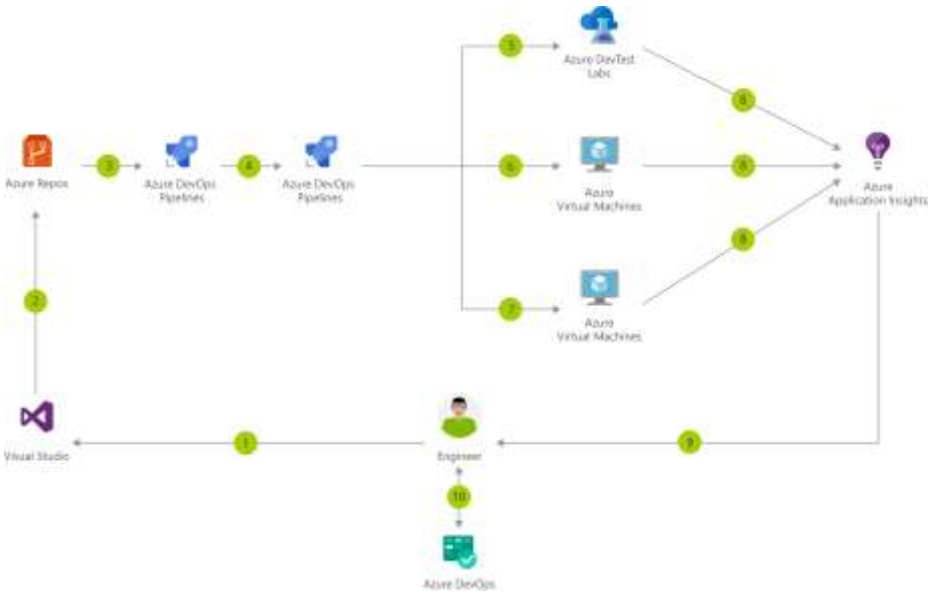

Figure 19 - Azure DevOps

Azure DevOps supports both public and private cloud configurations – the services include:

- Azure Boards – A work tracking system with Kanban boards, dashboards, and reporting
- Azure Pipelines – A CI/CD, testing, and deployment system that can connect to any Git repository
- Azure Repos – A cloud-hosted private Git repository service
- Azure Test Plans – A solution for tests and capturing data about defects
- Azure Artifacts – A hosting facility for Maven, npm, NuGet and Python packages

Each of these Azure DevOps services is open and extensible and can be used with all combinations of applications, regardless of the framework, platform or cloud. Built-in cloud-hosted agents are provided for Windows, Mac OS and Linux and workflows are enabled for native container support and Kubernetes deployment options, virtual machines, and serverless environments.

With all five services together, users can take advantage of an integrated suite that provides end-to-end DevOps functionalities. But, since they are broken into separate components, Azure DevOps gives users the flexibility to just pick which services to employ without the need to use the full suite.

Azure DevOps addresses the vendor lock-in problem from its early version by providing extensive integration with industry and community tools. With the many integrations

available, users can log in using SSO tools like Azure AD or communicate with their team via Slack integration while accessing both cloud and on-premises resources.

Azure Pipelines offers free CI/CD with unlimited minutes and 10 parallel jobs for every open-source project. As for Azure DevOps pricing, the basic plan for open source projects and small projects is free up to five users. For larger teams, the cost can range from $30 per month for 10 users to $90 per month for 20 users and so forth.

In summary, Azure DevOps is an all-in-one focused project tracking and planning tool mixed with Developer and DevOps tools for writing, building, and deploying code that's relatively quick and easy to use. But, while maintenance costs are decreased, developers only need an active subscription to have constant access to the latest version. Azure DevOps will indirectly utilize Azure Storage and computing services that will increase usage and impact costs.

Design Decision
ABC Company should use the newest features of Azure DevOps as one of the first transformational processes from IaaS to PaaS in Microsoft Azure. Benefits will be seen quickly by keeping provisioned infrastructure and applications in compliance.

VMWare VMs – Windows Server Assessment

The below represents the summary from the assessment created by the Cloud DB team at My IT Team. Parts of this assessment represent some of the highest memory and disk IOPS utilization on ABC Company systems. After migration, further refactoring to Azure Managed services can be performed in a phased manner to fully realize the benefits of Azure Managed Platform Services.

Windows Sizing Assumptions

This recommendation consists of an Azure VM replacing each existing VMWare/Hyper VM. High availability inherent in Azure through Azure Availability Sets exists without additional DR.

Assumption	Details
Target Location	East US 2 maps to Montvale and Toronto data centers
Target Storage disk	Managed Disks based on performance data of the disks

Assumption	Details
Azure Reserved VM Instances	MSOA Model consumption method was utilized in these cost estimations with 1-year reserved instances. Azure Migrate collector appliances were utilized.
Sizing criterion	VM Performance-based Duration: 1/1/2020 through 11/31/2020
Percentile Utilization	95th Percentile performance sample used for the right sizing
VM Series for rightsizing	DSv2_series, Dsv3_series, Dv2_series, Dv3_series, Ddv4_series, Dv4_series, Ddsv4_series, Dsv4_series, D_series, DS_series, Ev3_series, Esv3_series, Ev4_series, Edv4_series, Esv4_series, Edsv4_series, M_series, Fs_series, F_series, Fsv2_series, Lsv2_series
Comfort factor	2x buffer was utilized applied to CPU, RAM, disk and network data for VMs
Azure Hybrid Benefit	Software Assurance and eligibility applied to estimates
Azure pricing	As of 12/9/2020

Table 11 - Assessment Details

Sizes for virtual machines in Azure

This link references the available sizes and options for the Azure virtual machines you can use to run your apps and workloads. This Azure VM server assessment has run data collectors for approximately 4 weeks in ABC Company data centers to help with performance sizing recommendations. While the collection interval is 30 days, we have added a comfort factor to raise the confidence ratings to 5 stars.

Please review Azure server assessments to gain a better understanding of the sizing algorithms. The main assumption is to use a lift and shift strategy to migrate VMs from on-premises data centers to Microsoft Azure. Further refactoring assessments can be applied in later phases to ensure continuity through a phased design process.

Summary Table – Raleigh and Bakersfield

Subscription ID	xxx11xx-1x1x-11x1-xx1x-11x11xx111x1
Resource group	collector
Project name	MITCOLLECT011110project
Group name	Customer-Scoped-US and CA_Svrs
Assessment name	Cust_Scope_EastUS2_1yrRI_95_CF20_30d_AST
Assessment Type	Azure Server Assessment
Created on (UTC)	04/01/2025 6:10:16 PM
Total machines assessed	18
Machines not ready for Azure	0
Machines ready with conditions	0
Machines ready for Azure	18
Machines readiness unknown	0
Total monthly cost estimate USD	8182.13
Compute monthly cost USD	4689.12
Storage monthly cost USD	3493.01
Standard disks cost USD	3001.49
Standard SSD disks cost USD	0
Premium disks cost USD	491.52
Confidence rating	5-Stars

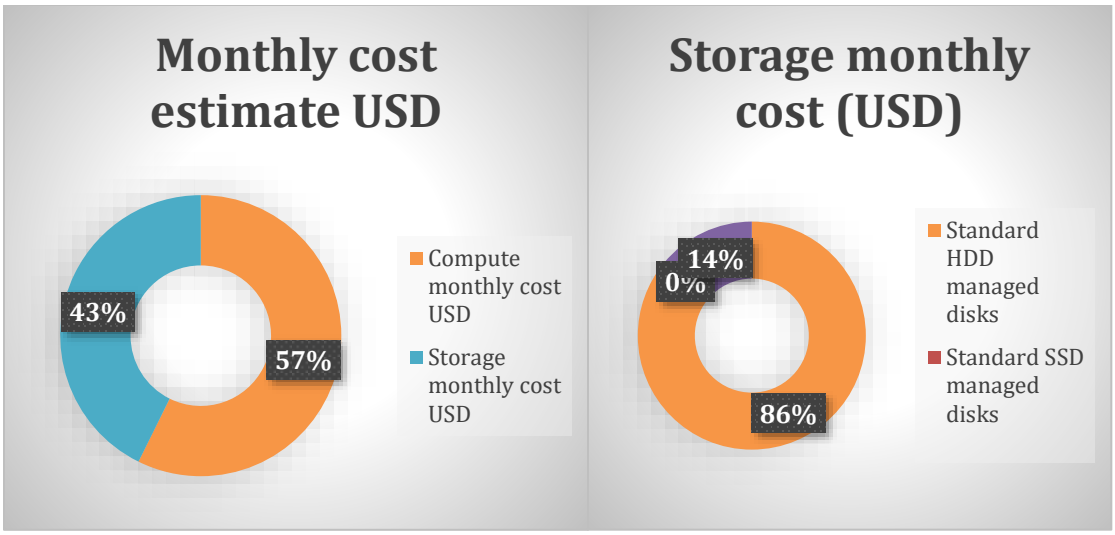

Figure 20 - Cost Estimates

Assessment List – Raleigh, NC

Machine	Category	Function	Apps Installed	Recommended size	Compute monthly cost USD	Storage monthly cost USD
LLMONDB04.hoya.local	SQL	SQL2014, VSTA3.0 RT, LAPS, endoPRO iQ	96	Standard_D4s_v3(40)	83.59	189.44 (P30, P20)
LLMONDB08.hoya.local	SQL	SP2013Foundation, SQL2014-2016, VSP2015, VSTA2015, MSBUILD Tools14, LAPS,endoPRO iQ	424	Standard_Ds4_v2(35)	485.81	28.55 (S10, S15,S15)
LLMONFTP01.hoya.local	App	LAPS, endoPRO iQ, IIS ADFS, ADRM, File Services	34	Standard_DS2_v2(35)	54.08	276.48 (P40, S30)
LLMONDB05.hoya.local	SQL	SQL2014, VSTA3.0 RT, LAPS,endoPRO iQ	89	Standard_Ds4_v2(35)	485.81	38.98 (S10, S15, S20)
LLMONSP04.hoya.local	App	SP2013, Project2013, SSMS, VS2015, MSBUILDTools14, VSTA2015, Workflow Server, endoPRO iQ	96	Standard_F16s_v2(41)	291.92	66.56 (P20)
LLMONDB07.hoya.local	SQL	SQL2014-2016RTM, VSP2015, VSTA2015, MSBUILD Tools14, LAPS, SSMS-AR	353	Standard_Ds4_v2(35)	485.81	38.98 (S10, S15, S20)

Table 12 - Raleigh Site

Assessment List – Bakersfield, CA

Machine	Category	Function	Apps Installed	Recommended size	Compute monthly cost USD	Storage monthly cost USD
SQL12CRM15	SQL	SQL2012SP4, VSTA2017, SSMSA, LAPS	25	Standard_D15_v2(29)	772.75	327.68 (S60)
SQL12SRV1	SQL	SQL2012SP4, VSTA3.0 RT, MSOFFICE2010, SSDT, LAPS	61	Standard_D13_v2(29)	309	204.80 (S30, S50)
SQL12SRV2	SQL	SQL2012RTM, VSTA2015, VSE2019, VSTA3.0 RT, SSDT, LAPS, TFS2015	117	Standard_D13_v2(29)	309	204.80 (S30,S50)

Machine	Category	Function	Apps Installed	Recommended size	Compute monthly cost USD	Storage monthly cost USD
SQLSRV (Windows 2003 Standard)	SQL	SQL2000SP4, Infor PM BPA, OpenEdge10 ABL	52	Standard_F4s_v2 (41)	73	33.99 (S15x3)
PROPHIX01.hoya.local	App	SQL2012, VSTA3.0 RT, LAPS, OpenEdge10, ProPhix11	43	Standard_D8_v4 (41)	165.42	40.96 (S30)
scibe-srv.hoya.local	App	OPP2010, LAPS, Scribe Insight with Adapters for CRM and ODBC	26	Standard_DS2_v2(35)	54.08	21.76 (S20)
LXARGPORTAL0001.hoya.local	App	VSE2019, LAPS, SQL2016, ADXSytudioPortals for MS Dynamics, Azure Emulators	31	Standard_F16s_v2 (41)	291.92	43.52 (S20x2)
LXARGCRM0002.hoya.local	App	Dynamics CRM2015, LAPS	15	Standard_D16s_v4 (41)	330.75	81.92 (S30x2)
LXARGHYPEPMA	App	EPM Automate, TeamViewer	5	Standard_D4_v4 (41)	82.67	21.76 (S20)
LXARGFILE0001.hoya.local	Shared Services	LAPS, MergeModule2012	30	Standard_D8_v4 (41)	165.42	1474.56 (S50, S70x2)
LXARGPRINT0001.hoya.local	Shared Services	LAPS, MergeModule2012	30	Standard_D4_v4 (41)	82.67	40.96 (S30)
SX-RDP.hoya.local	App	LAPS, TightVNC, MergeModule2012	31	Standard_D8_v4 (41)	165.42	22.66 (S15x2)

Table 13 - Bakersfield Site

For additional details such as the disk sizing, please see attached excel sheet.

Cust_Scope_EastUS 2_1yrRI_95_CF20_30(

Design Decision
As these Virtual Machines will be used for long term, we recommend reserving the instances for minimum 1 year to save cost**s** up to 41% as of this writing.

ABC Company Group has workloads installed in two datacenters across the enterprise. In addition, the single windows VM can contain multiple applications as well as a single distributed application can be installed on multiple VMs.

Workload Relationship to Infrastructure – Raleigh & Bakersfield

The workload relationship to Infrastructure or all VMs is located on the 2nd sheet of the embedded Application Discovery document in the next section:

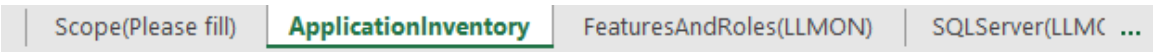

Application Discovery

The following excel sheet provides a list of VMs with applications discovered:

PAMC_Apps_LLMO NCOLLECT01_Scope.

Application Dependencies

Dependency analysis identifies dependencies between discovered on-premises virtual machines and hosts. It provides the following information:

- You can identify machines that must be migrated together. This is especially useful if you're not sure which machines are part of an app deployment that you want to migrate to Azure.
- You can identify whether machines are in use, and which machines can be decommissioned instead of migrated.
- Analyzing dependencies helps ensure that nothing is left behind and thus avoids surprise outages after migration.

After Discovery, Dependency data polling begins:

- The Azure Migrate appliance polls TCP connection data from machines every five minutes to gather data.
- Data is collected from guest VMs via vCenter Server, using vSphere APIs.
- Polling gathers the following data points:

  - Name of processes that have active connections.
  - Name of application that run processes that have active connections.
  - Destination port on the active connections.
- The gathered data is processed on the Azure Migrate appliance, to deduce identity information, and is sent to Azure Migrate every six hours.

The following excel sheet provides information about the application dependency data for both data centers – Raleigh and Bakersfield.

PAMC_PCI_AppDMapping_All.xlsx

The best way to read the results is to utilize the pivot tables on the following named tabs:

- **PAMC-Pivot1-BySourceSvr** – Provides all TCP connections from Source Server to Destination Server on destination listening TCP port.
- **PAMC-Pivot1-ByDestPort –** Provides all TCP connections by Destination Port of Destination Server
- **PCI-Pivot1-BySourceSvr -** Provides all TCP connections from Source Server to Destination Server on destination listening TCP port.
- **PCI-Pivot1-ByDestPort -** Provides all TCP connections by Destination Port of Destination Server

For example, the following table provides information on which servers the SharePoint Server talks to on a regular basis:

Source IP	(All)		
Source server name	Destination server name	Destination po	Time slot
(blank)			
External			
PMTEST			
PMBCKUP-SAP			
LLMONDB07.hoya.local			
LLMONTFSTEMP01.hoya.local			
LLMONBUVE01			
LLMONOWA01.hoya.local			
LLMONAS01.hoya.local			
LLMONEIQ02.hoya.local			
LLMONDB08.hoya.local			
LLMONSP04.hoya.local			
LLMONSP04.hoya.loc	(blank)		
LLMONSP04.hoya.loc	External		
LLMONSP04.hoya.loc	LLMONSE012-RODC		
LLMONSP04.hoya.loc	LLMONSE012-RODC	88	
LLMONSP04.hoya.loc	LLMONSE012-RODC	135	
LLMONSP04.hoya.loc	LLMONSE012-RODC	389	
LLMONSP04.hoya.loc	LLMONSE012-RODC	445	
LLMONSP04.hoya.loc	LLMONSE012-RODC	636	
LLMONSP04.hoya.loc	LLMONSE012-RODC	3268	
LLMONSP04.hoya.loc	LLMONSE012-RODC	49155	
LLMONSP04.hoya.loc	LLMONSP04.hoya.local		
LLMONSP04.hoya.loc	LLMONSP04.hoya.local	443	
LLMONSP04.hoya.loc	LLMONSP04.hoya.local	9355	
LLMONSP04.hoya.loc	LLMONSP04.hoya.local	9355	1/7/2021 12:00-18:00
LLMONSP04.hoya.loc	LLMONSP04.hoya.local	9355	1/27/2021 18:00-0:00
LLMONSP04.hoya.loc	LLMONOWA01.hoya.local		
LLMONSP04.hoya.loc	LLMONOWA01.hoya.local	80	
LLMONTMG001.hoya.local			
PMSHIP			

Figure 21 - Server Dependency

Implementation and Migration Strategy

Using the table of contents as an implementation guide (shown below), configure the Governance Model, Network Architecture with Security and define the storage. Next, configure Backups and Disaster Recovery before setting up the Hybrid identities. Throughout the process, it is important to reference the Cloud Adoption Framework for more detailed online guidance.

- Configure Governance Model
  - Enrollment Model
  - Management Groups
  - Naming Standards
  - Resource Groups
  - Tagging Standards
  - Azure Policy
  - Resource Locks

- Setup Network Architecture
  - Network
  - VNET
  - SUBNET
  - VNET Peering
  - IP Schema
  - Network Topology
  - Load Balancers
  - Azure Connectivity
- Establish Security
  - Network Security Group
  - Azure Firewall
  - DDoS Protection
  - Azure Information Protection (Data Security)
  - Storage Account – Data in Transit
  - Azure Advisor
  - Azure Sentinel (SIEM)
  - Azure Security Center (End Point Protection)
- Plan for Azure Virtual Machine Workloads
  - Virtual Machines (VMs)
  - VM Families and Tiers
  - VM Management and Monitoring
  - High Availability
- Define Storage
  - Virtual Machine Storage
  - Management of storage accounts
  - Azure Storage Encryption
- Setup Disaster Recovery
  - Backups
  - Disaster Recovery Standards
  - Backup Considerations
  - Azure Site Recovery
  - Recovery Plans
  - Azure DNS
- Configure Identity Management (Hybrid)
  - Azure Active Directory & Domain Services
  - User Management
  - Role Based Access Control (RBAC)
- Setup Management & Operations
  - Infra Monitoring

- Network Monitoring
- Automation
- Microsoft Azure DevOps