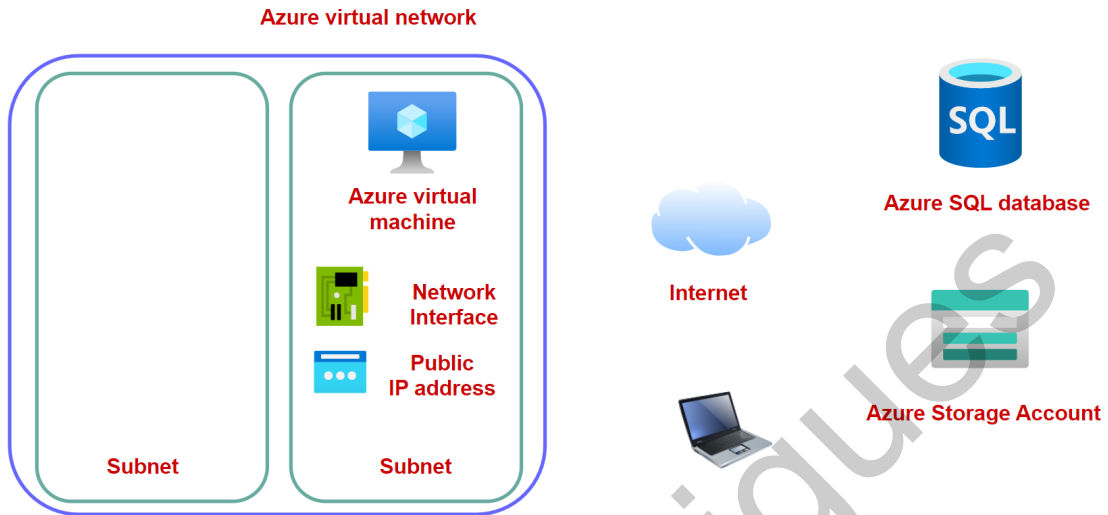
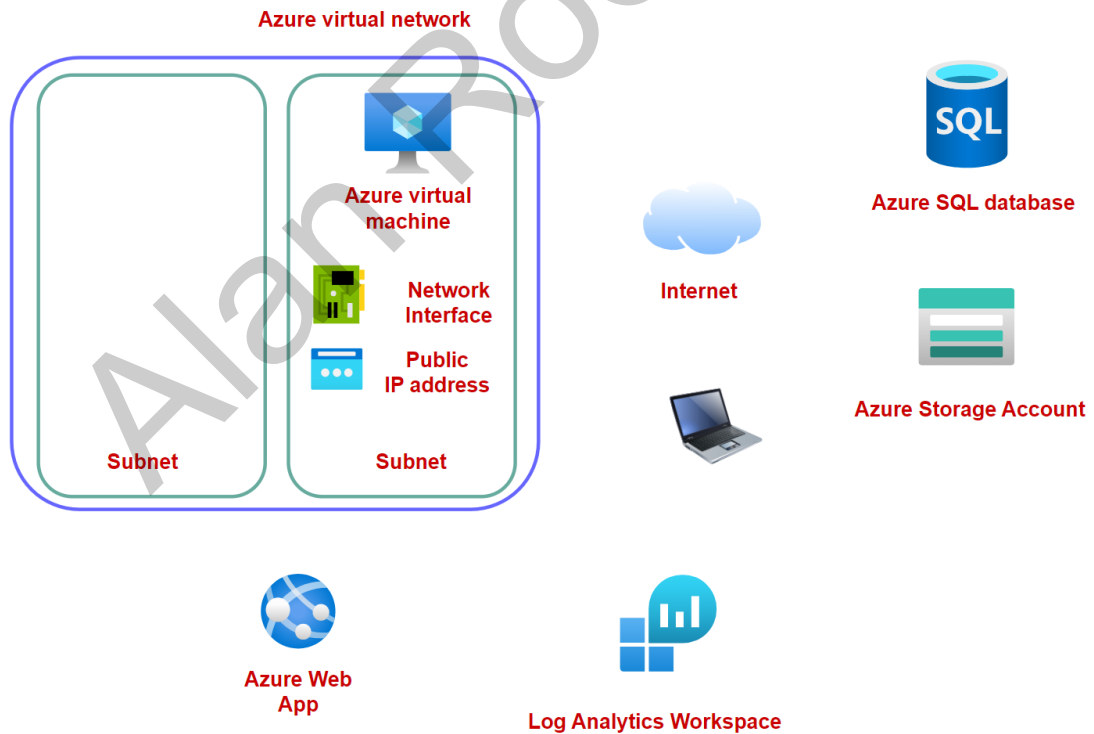


Design Monitoring

Architecture

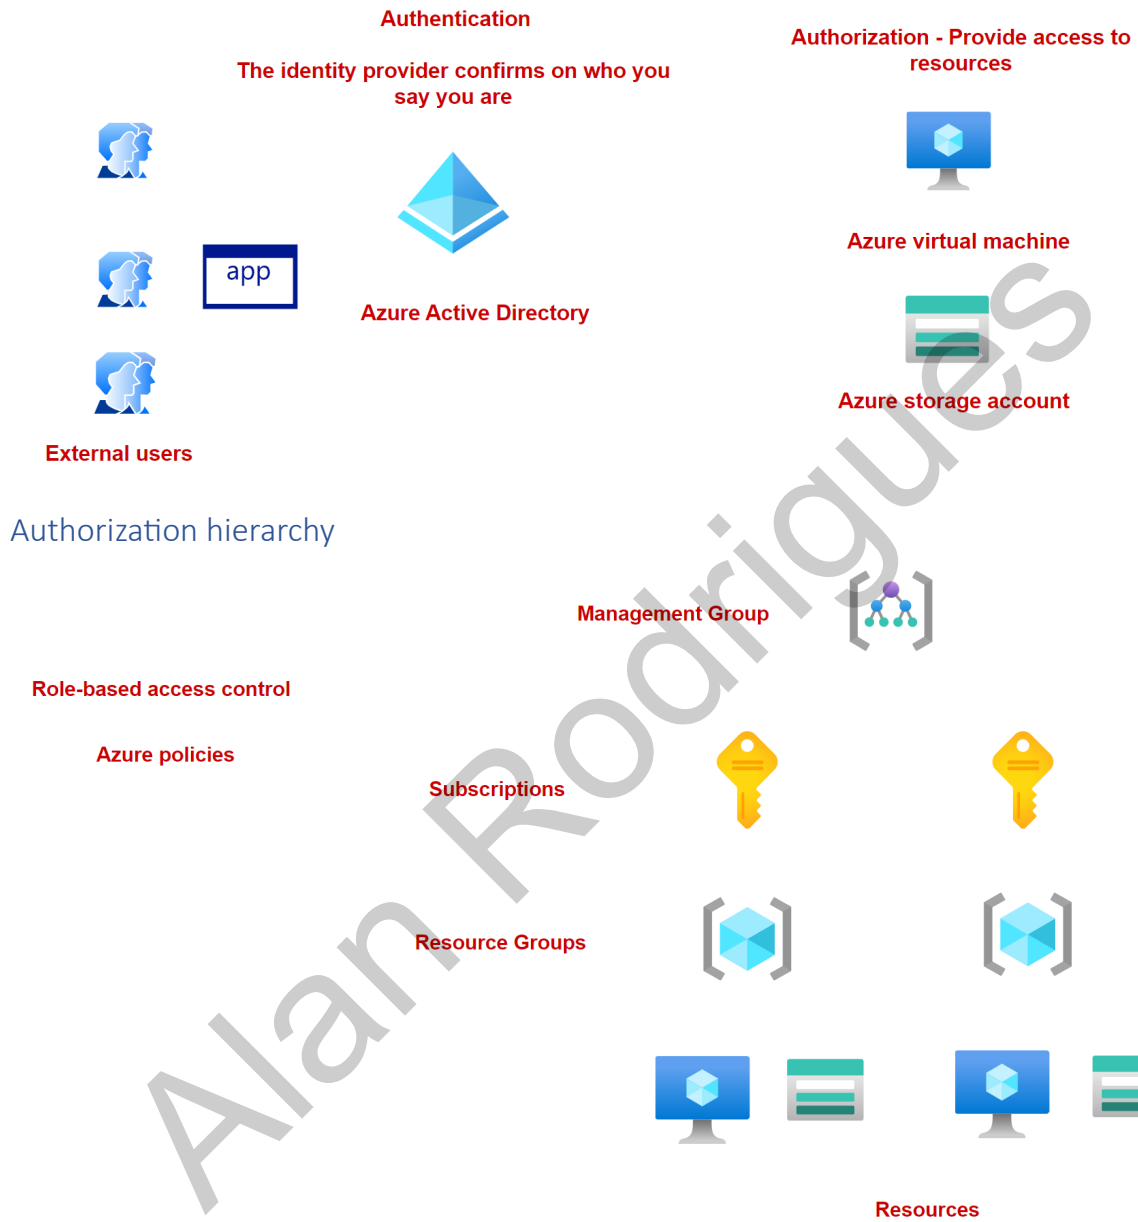


Review- Monitoring- Log Analytics workspace- Azure VM's



Design Identity and Security

Authentication and Authorization



What are Azure AD Access Reviews



Azure Blueprints



Azure Blueprints- Quick Note



Azure Blueprints

Orchestrate the deployment of artifacts to Azure

All subscriptions under a Management Group must trust the same Azure AD tenant



default directory

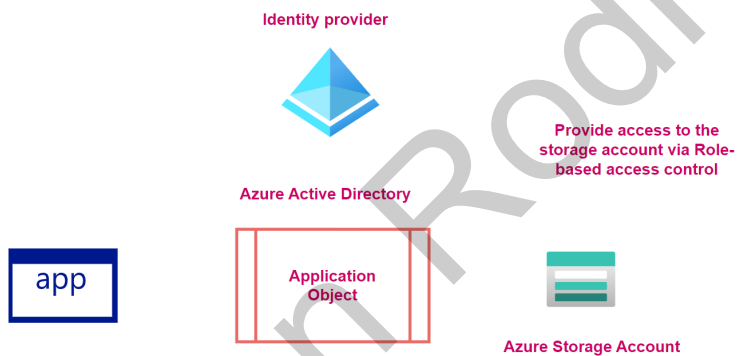


cloud2hub

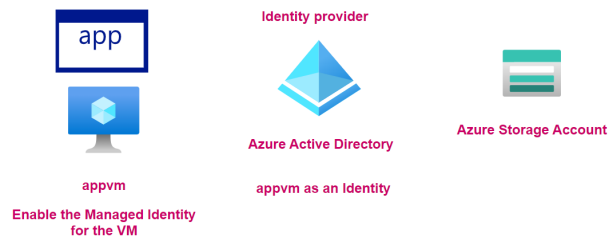
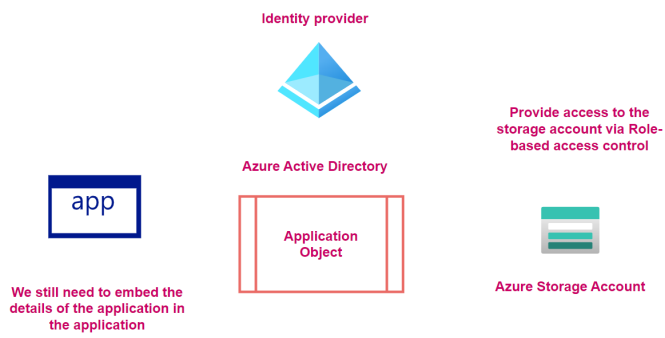
Here we need two management groups

We would need two Blueprint definitions that can be assigned to each Management Group separately

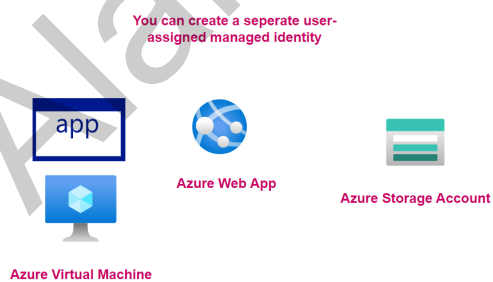
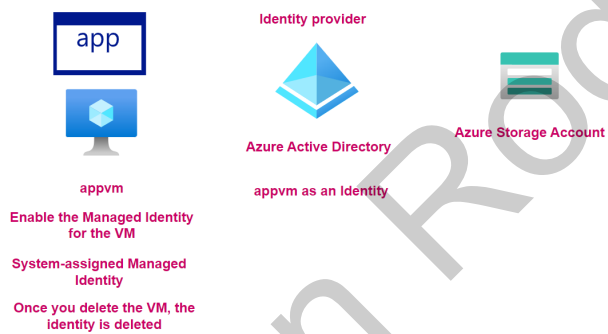
Azure AD- Application Objects



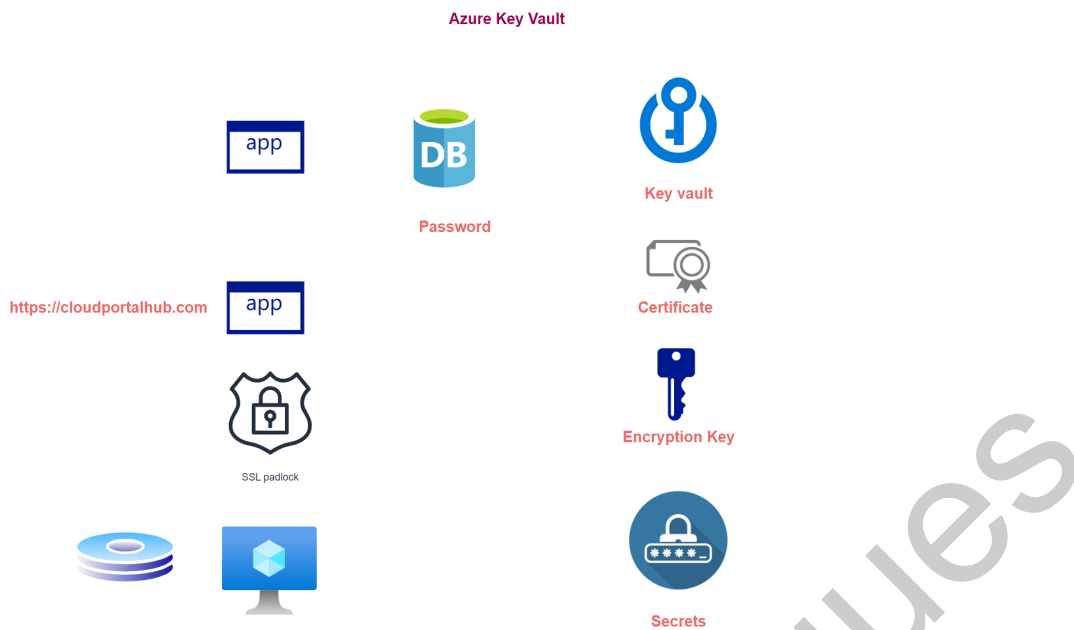
Managed Service Identity



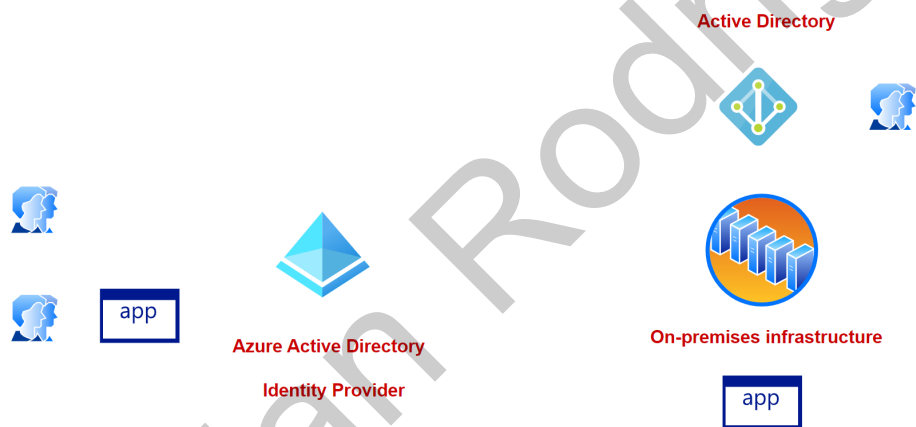
User assigned managed identities



Overview of the Azure Key vault



Azure AD Application Proxy



You can have connections between the on-premises environment and Azure AD when it comes to identities

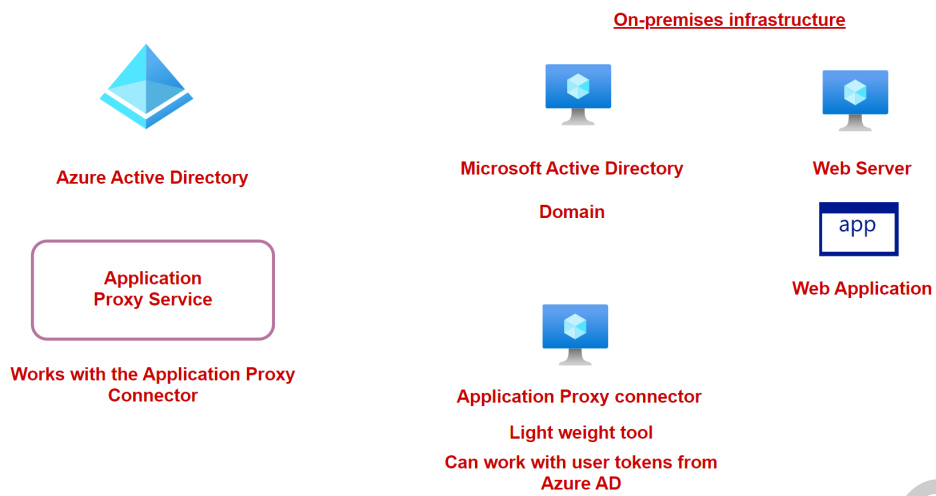
Requirement

We want our users defined in Azure AD to be able to access the web application hosted on our on-premises servers.

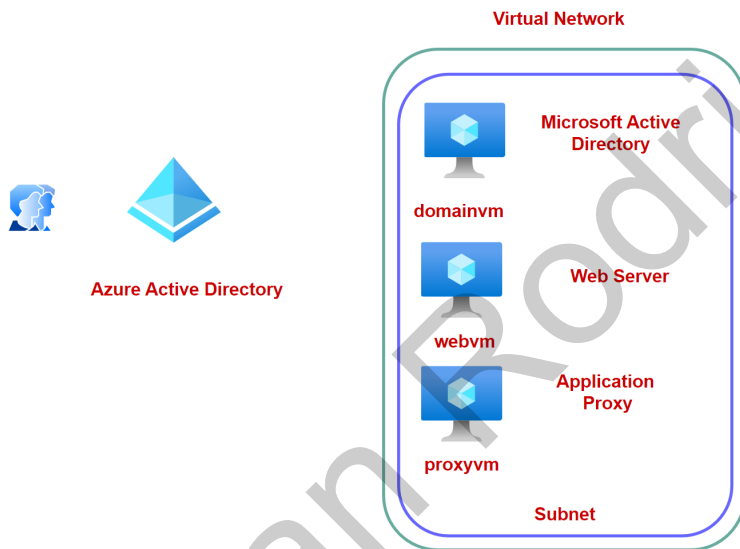
Earlier one step in the implementation would be to expose the application onto the Internet.

But that would add a security risk wherein you are directly exposing the server onto the Internet.

We could instead make use of Azure Application Proxy



Architecture



Review on the different Authentication processes



Azure Active Directory



Users log in via a user name and password

Additional security layer of Multi-Factor Authentication



Application gets access to resources via an Application object



Users log in via an Application. Here Azure AD authenticates the user

Our sample application was actually making use of OAuth

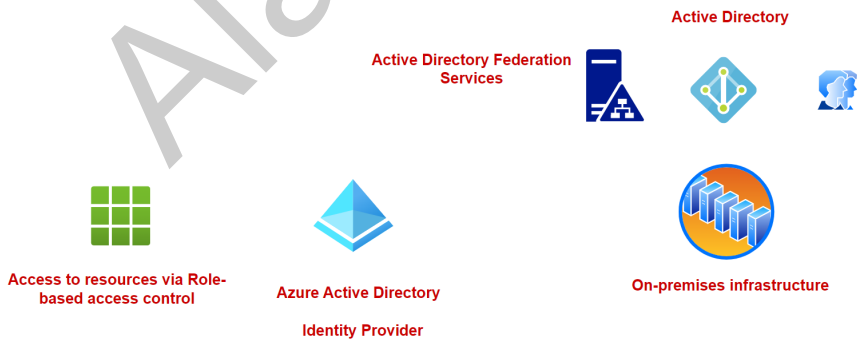
The application calls an OAuth endpoint on Azure and gets a token.

The token will have the permissions the user has when it comes to authorization to Azure resources.



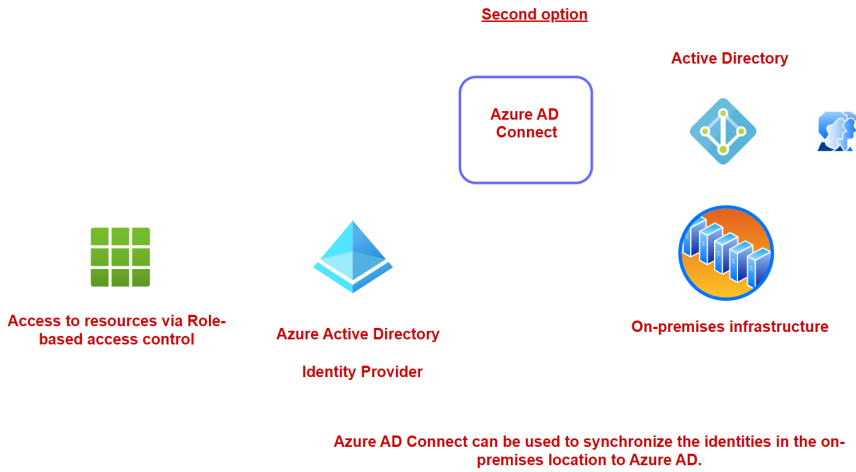
Here the company wants their on-premises identities to have access to Azure-based resources.

One option

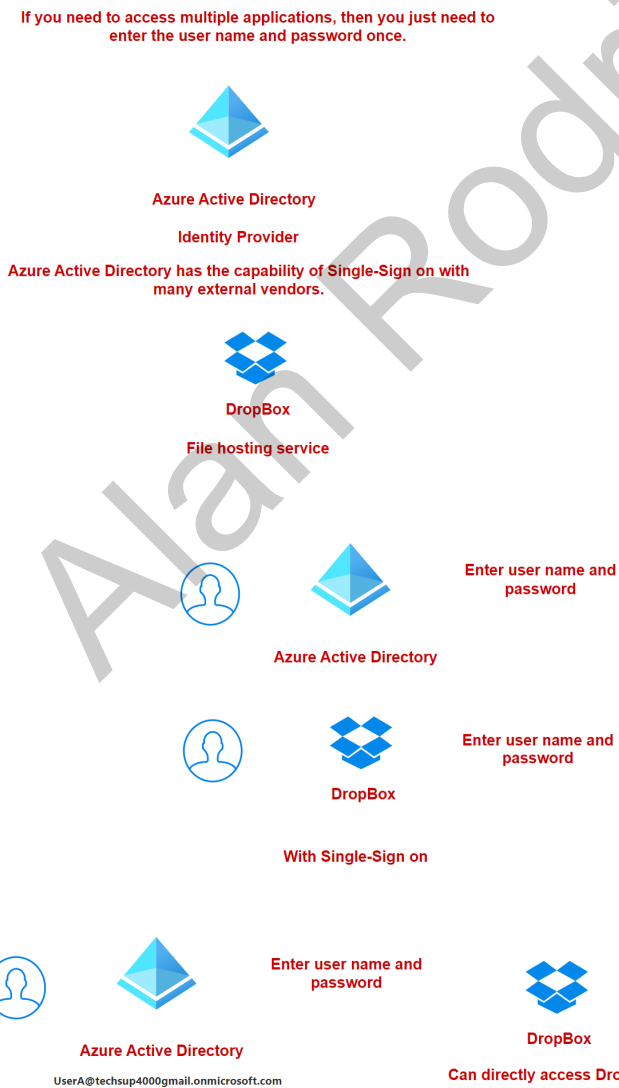


Active Directory Federation Services can make use of the user identities that have already authenticated themselves via the on-premises Active Directory.

Active Directory Federation Services can then work with Azure AD to authorize the use of Azure resources to the identities.



Enterprise Applications- Single Sign On



Azure AD needs to support Single-Sign on with the external application

DropBox needs to support Single-Sign on with Azure AD

And you need to configure Single-Sign on in both areas

Design Data Storage

Different data stores overview

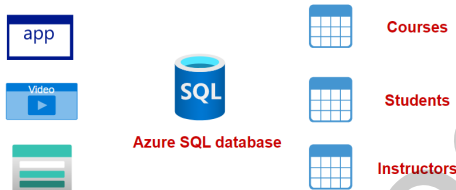
Different data stores



Application serves videos to users

We need a data store for the videos

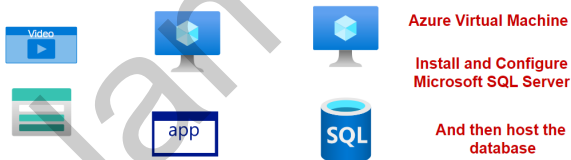
We need a data store for storing information with regards to students, instructors, courses etc.



Azure Storage Account

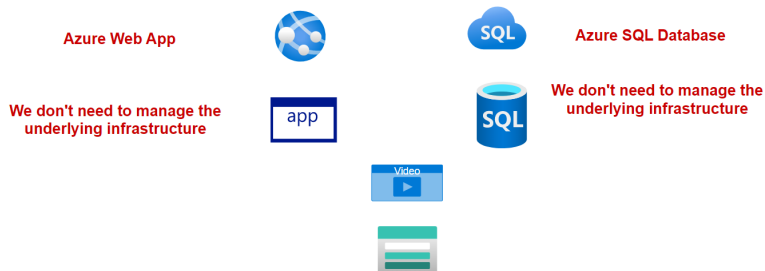
Have relationships between the tables

When it comes to using Microsoft SQL Server, we have a couple of options



We could even host our web application on an Azure virtual machine

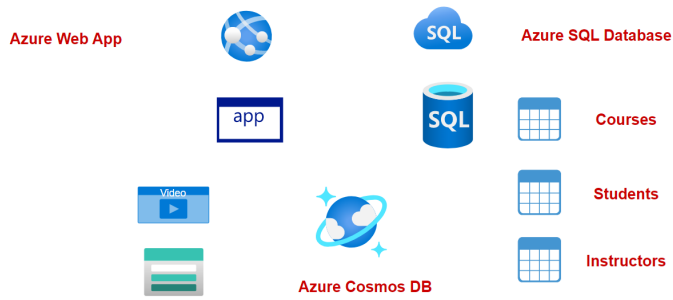
Or we can use the PaaS services - Platform as a service



We don't need to manage the underlying infrastructure

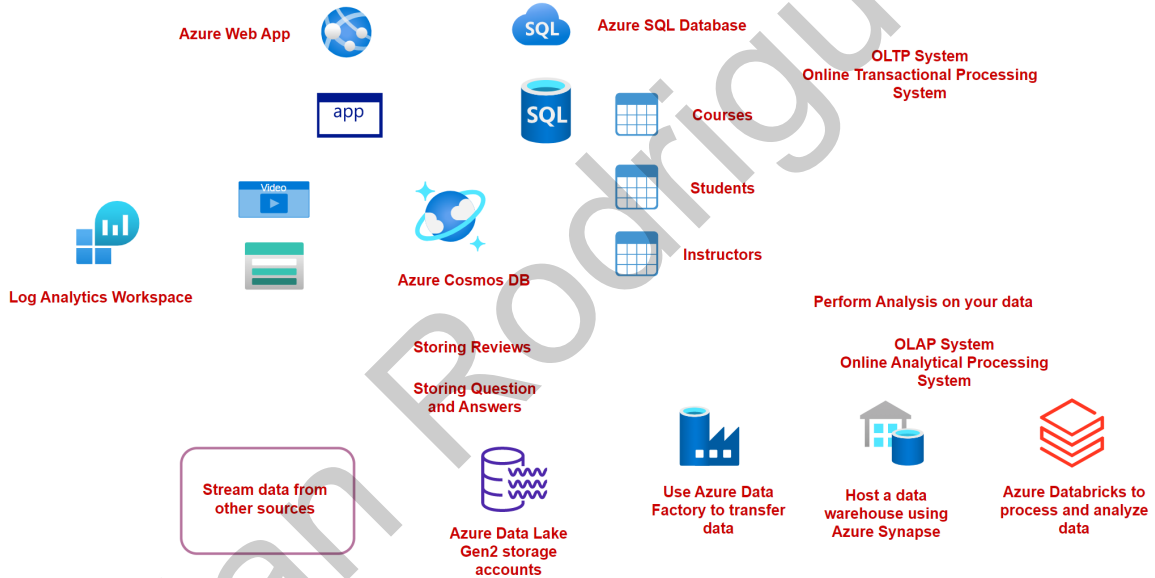
We don't need to manage the underlying infrastructure

Using Azure Cosmos DB - NoSQL data store

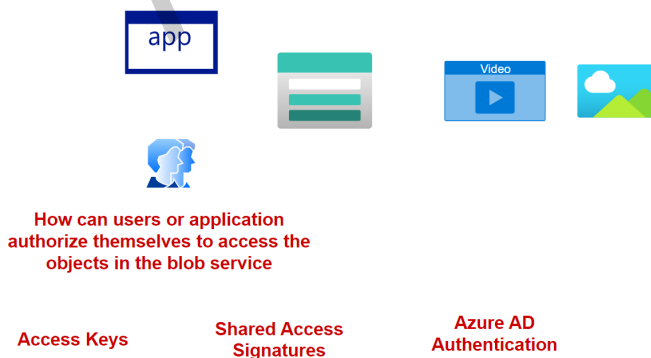


Storing Reviews

Storing Question and Answers



Azure Storage Accounts- Different authorization techniques – Review



Azure Blob Storage tiering

Azure storage access tiers



Objects

There is a cost for storing objects

There is a cost for accessing objects

Companies might store millions of objects in a storage account

Use case - Initial there could be some objects that are accessed quite frequently. Then after some time, maybe a week or two, those objects are accessed less frequently.

Can a company save on costs when it comes to less frequently accessed objects.



An object can be set to a particular tier

Hot
Access
tier

This is optimized for objects accessed more frequently
Here you have high storage costs and lower access costs

Here the data needs to be stored for at least 30 days

Cool
Access
tier

This is optimized for objects accessed or modified infrequently
Here you have lower storage costs but higher access costs when compared with the Hot access tier.

Here the data needs to be stored for at least 180 days

Archive
Access
tier

This is optimized for objects that are rarely accessed
Here you have lower storage costs but higher access costs when compared with the Cool access tier.

Good for long-term backups.

You can set the Hot and the Cool access tier at the storage account level.

You can set the Hot, Cool and Archive access tier at the blob level.

Premium Storage Accounts

Premium storage accounts



Premium block blobs

This is used when you need high performance when it comes to storage and access to data.

Here the data in the background is stored on solid-state drives. These are optimized for low latency.

Here the file transfer is also much faster.

Workloads - Streaming , Machine Learning

You have higher storage costs but lower transaction costs

Data redundancy

Performance ⓘ *

Premium account type ⓘ *

Redundancy ⓘ *

Locally-redundant storage (LRS):
Lowest-cost option with basic protection against server rack and drive failures. Recommended for non-critical scenarios.

Zone-redundant storage (ZRS):
Intermediate option with protection against datacenter-level failures. Recommended for high availability scenarios.

Locally-redundant storage (LRS) ▼

You can't set the access tiers

Premium file shares

Here again you get high performance and low latency

Backed by solid-state drives for storage

Data redundancy

Performance ⓘ *

Premium account type ⓘ *

Redundancy ⓘ *

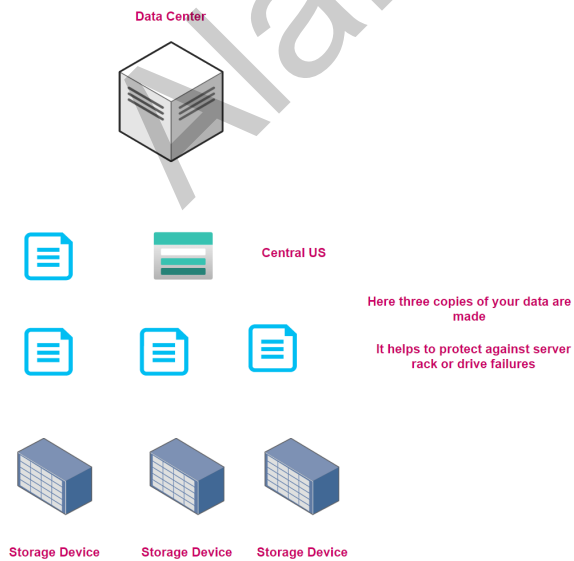
Locally-redundant storage (LRS):
Lowest-cost option with basic protection against server rack and drive failures. Recommended for non-critical scenarios.

Zone-redundant storage (ZRS):
Intermediate option with protection against datacenter-level failures. Recommended for high availability scenarios.

Locally-redundant storage (LRS) ▼

Azure Storage Accounts- Data Redundancy

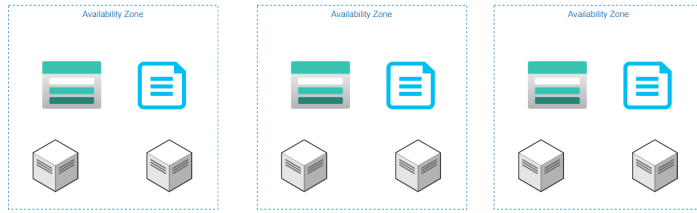
Locally-redundant storage



Zone-redundant storage

This helps to protect against data center level failures

Here data is replicated synchronously across three Azure availability zones

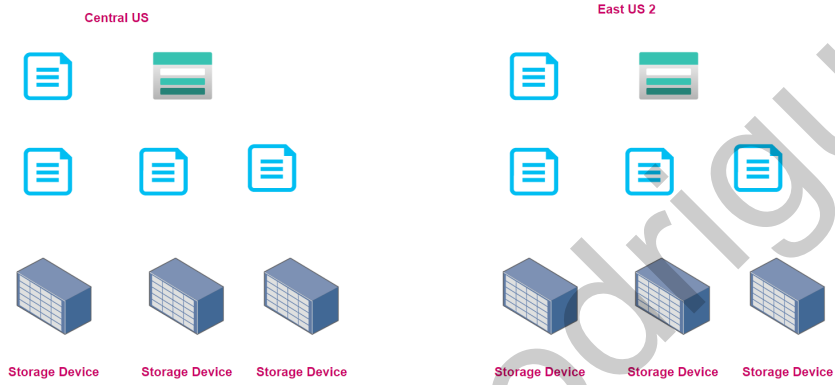


Central US

Each availability zone is a separate physical location with independent power, cooling and networking

Geo-redundant storage

Here data is replicated to another region

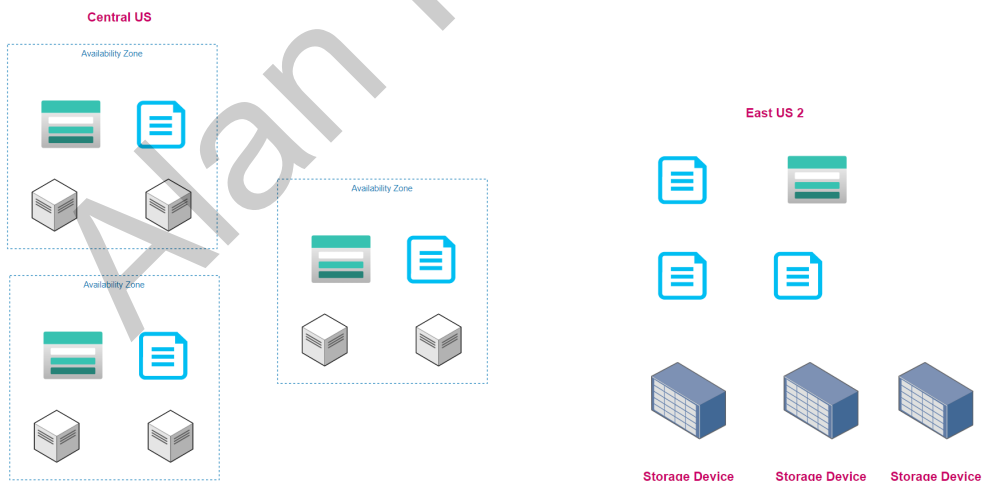


Data is copied three times in the primary region using LRS

Data is copied three times in the secondary region using LRS

Geo-zone-redundant storage

Read Access geo-zone-redundant storage



Lab- Azure Cosmos DB

Azure Cosmos DB



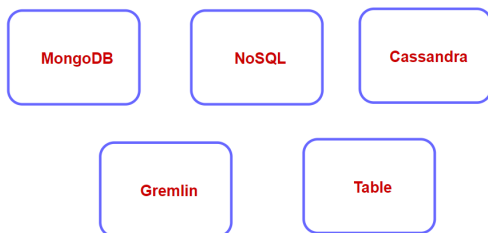
Fully managed NoSQL database system

Here the infrastructure is managed for you

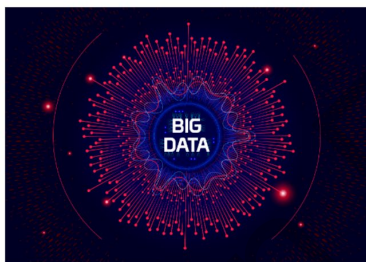
You get fast access to your data

You can also create multiple read and write replica's of your data to multiple locations

You can also choose from different API's



Lab- Azure Data Lake Storage Gen 2



Working with large data sets

Here the data arrives in large volumes

The data arrives at a fast rate

Azure Data Lake Storage Gen2



This service is built on top of Azure Blob storage

Gives the ability to host an enterprise data lake on Azure

You also get the feature of a hierarchical namespace on top of Azure Blob storage

A data lake is used to store large amounts of data in its native, raw format

Data lakes are optimized for storing terabytes and petabytes of data

The data could come from a variety of data sources

The data itself could be in various formats - Structured, semi-structured and unstructured data

What is Azure Synapse Analytics

Data Warehousing

Centralized repository of data taken from different data sources

Normally used by business for analyzing data



Let's say senior management wants to know

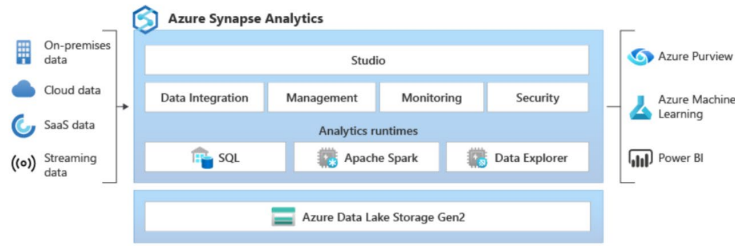
1. Purchases done per month per region
2. Most popular courses for the month per region
3. Times during the day where purchases are made the most



1. SQL Data Warehouse would again be powered by some database engine
2. Your data would be structured in the form of tables
3. But the data warehouse would store larger amounts of data
4. Data in the data warehouse is then analyzed accordingly
5. You need a lot of processor power to analyze the data
6. Just like a SQL database, you can fire SQL queries against a SQL data warehouse
7. Business users can also use reporting tools to visualize the data in the warehouse

Modern data warehouse architecture





Reference - <https://learn.microsoft.com/en-us/azure/synapse-analytics/overview-what-is>

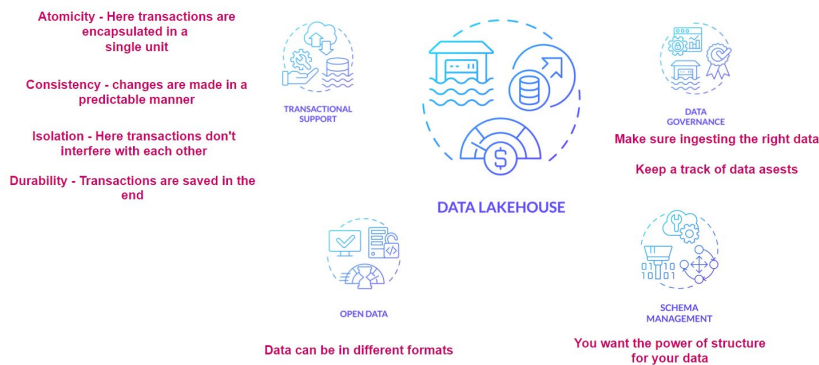
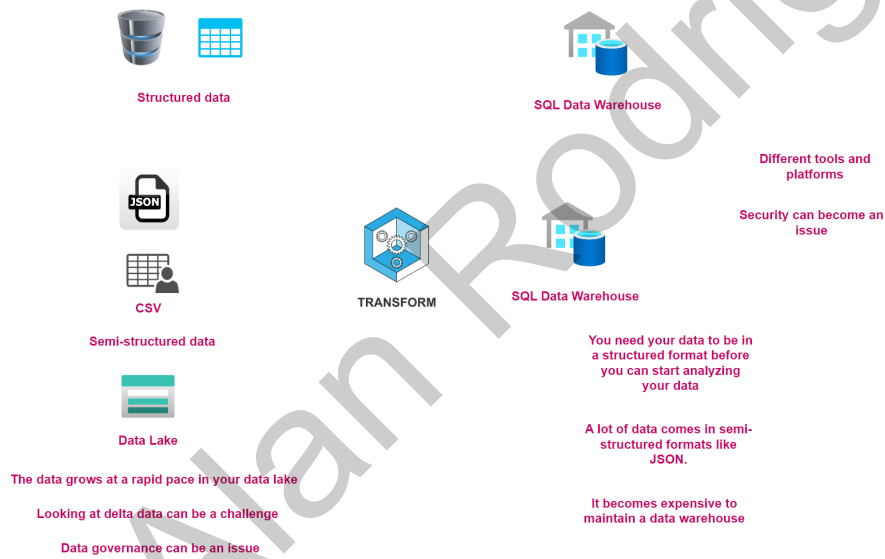
What is Azure Databricks

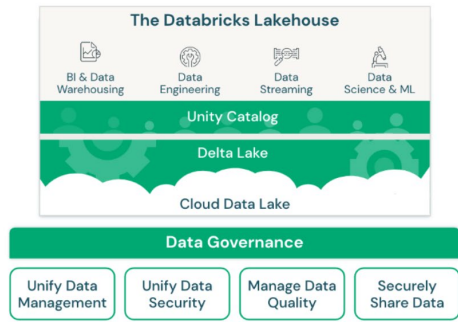
Databricks

This is a company that provides a platform for developing and maintaining enterprise-grade data solutions.

The creators of Spark went ahead and created the Databricks platform

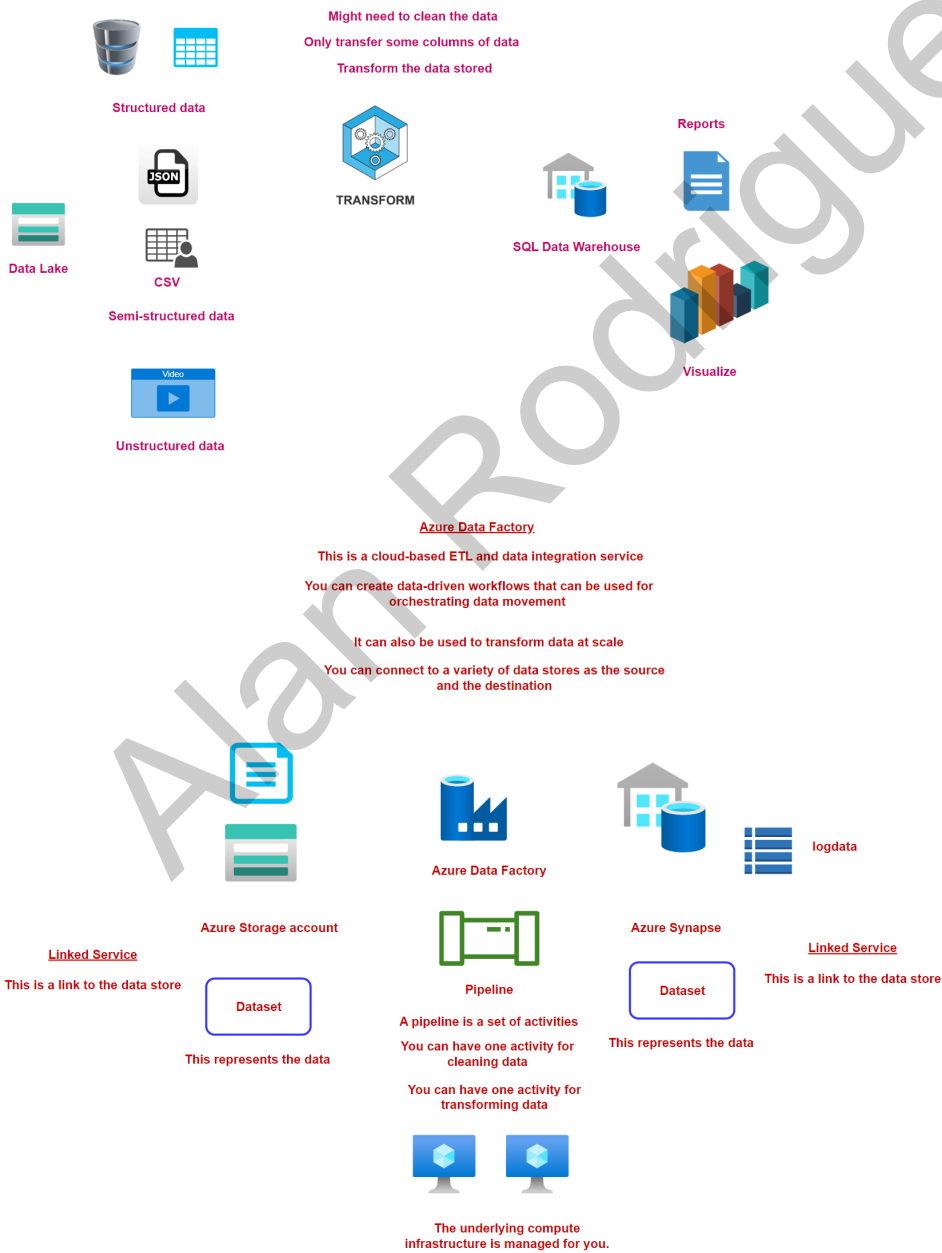
Ingest data Process data Work with SQL Data Exploration
Machine Learning Security and governance Visualizations





<https://docs.databricks.com/lakehouse-architecture/data-governance/index.html>

What is Azure Data Factory

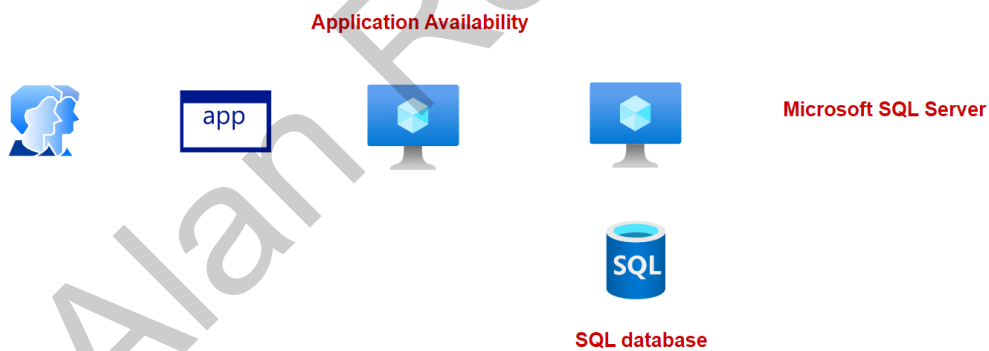


Azure Data Factory- Self-hosted runtime



Design Business Continuity

Review on Application Availability



What happens if the Virtual Machine hosting the application goes down?

What happens if the Virtual Machine hosting the SQL database goes down?

Do you have backup strategies in place?

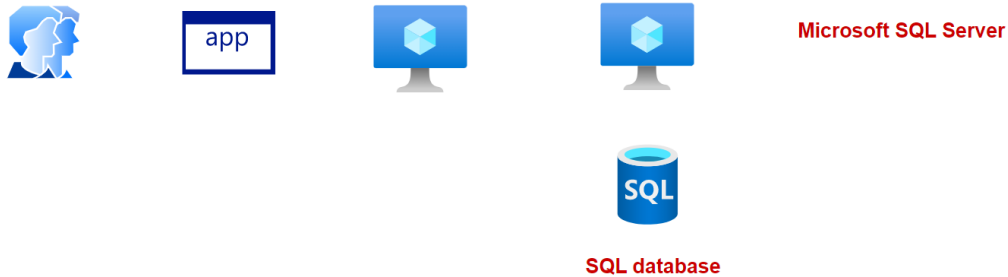
How much revenue does the company lose when an outage occurs?

Business Users losing confidence.

Do you need to look at disaster recovery scenarios, an entire region going down.

Recovery Point and Recovery Time Objective

Application Availability



RPO - Recovery Point Objective

What is the maximum data loss that is acceptable by business after a failure

Let's say that the Virtual Machine hosting the database goes down at 9:00 a.m

You only have backup's of data till 8:30 a.m

This means that you will lose 30 mins of data.

Is this acceptable by business.

It's impossible to have zero outages and also have zero loss of data.

Even using tools to minimize loss comes at a price.

It's all about creating a good business case on business loss vs cost of using different recovery services.

RTO - Recovery Time Objective

How long does it take for the entire infrastructure to come back online after a failure.

Again this is dependent on business requirements and how much loss the business can take

And again there is a cost for using tools that can decrease the Recovery Time objective

Also you need to consider the RTO of your entire system.

If it takes 10 minutes to get the application compute infrastructure up and running after an outage but it takes 30 minutes for the SQL database Infrastructure to get up and running, your final RTO will be 30 minutes.

Concepts

For every SQL database , the engine maintains a transaction log

This records all transactions and database modifications that are made by the transactions.

This allows you to bring back the database to a consistent state if there is a failure.

A SQL Server database will also have two files at the least.

Data file - This contains the data and objects when it comes to tables, indexes, stored procedures and views.

Log file - This is used to log the database transactions



The data in the end is written to files

.mdf - Data files

.ldf - Log files

To protect against failure we need to protect the disks that store the data

We also need to keep the compute infrastructure highly available

Decouple the availability of the compute and data infrastructure

Azure SQL database backup



Full database backup - Here the entire database that includes the transaction logs. This backup is taken once a week.

Differential backup - This includes the data that has changed since the last full database backup. This is taken every 12 to 24 hours.

Transaction log backups - log of the transactions. This is taken every 10 minutes.

With this automated backup, you can perform a point in time restore of the database.

Because what happens if someone deletes records by mistake in the database and you want to recover the data.

So you can restore the database using the backups. You can create a new database on the same server.

You can also restore a deleted database, but this needs to be done on the same server.

Recovery Point Objective - RPO - 10 minutes based on the compute size and the database activity.

Recovery Time Objective - RTO - Can take hours depending on the database size.

By default the backups are retained for 7 days.

Except for the Basic DTU tier databases, you can extend the retention to 35 days for the backups.

Azure SQL database- Backup Redundancy

SQL database backup storage redundancy



First lets understand the architecture

Service tier	High availability model	Locally-redundant availability	Zone-redundant availability
General purpose (vCore)	Remote storage	Yes	Yes
Business Critical (vCore)	Local storage	Yes	Yes
Hyperscale (vCore)	Hyperscale	Yes	Yes
Basic (DTU)	Remote storage	Yes	No
Standard (DTU)	Remote storage	Yes	No
Premium (DTU)	Local storage	Yes	Yes

<https://learn.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla?view=azuresql&tabs=azure-powershell>

As an example lets consider the Standard (DTU) Model

Has the transient and cache data on local SSD (Solid State Drives)



Seperate compute layer that runs the database engine



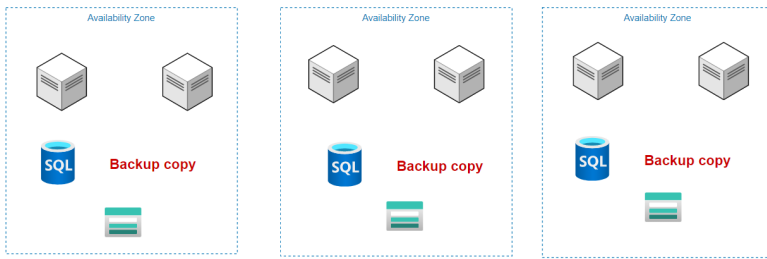
The data is stored in Premium storage and is backed by LRS

The actual data is stored in Azure Storage Accounts

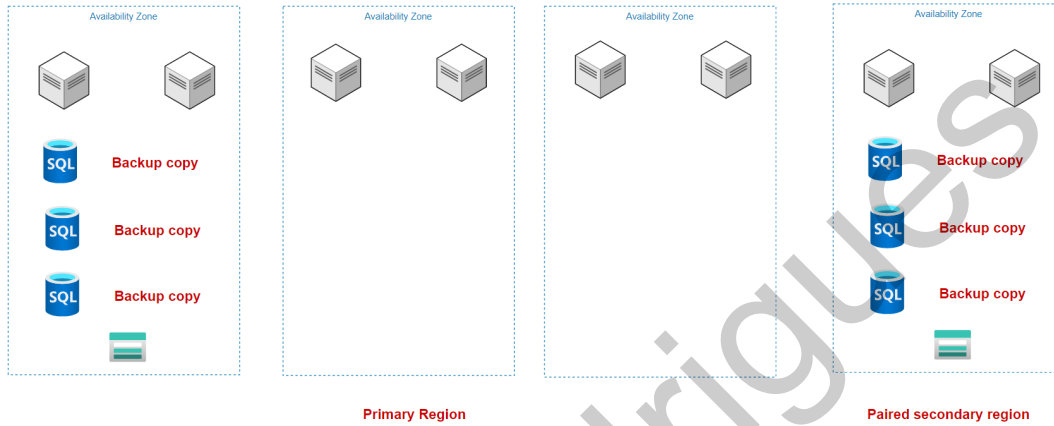
If anything happens to the compute layer, another compute machine will be created that will connect to the data in Azure Storage

When you choose the Premium Tier you can also choose to make the database itself zone-redundant

Zone-Redundant storage

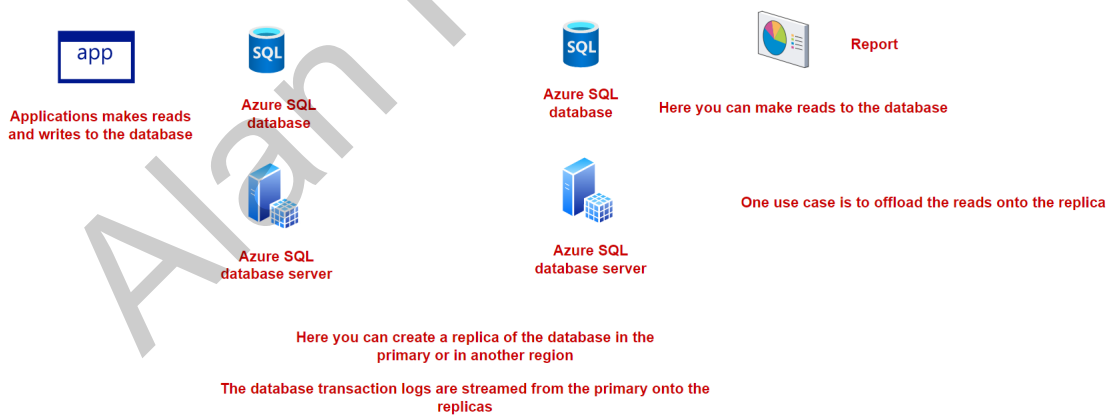


Geo-Redundant storage



Azure SQL database- Active geo-replication

Azure SQL database - Active geo-replication



Another use-case



Azure SQL
database



Azure SQL
database



Azure SQL
database server



Azure SQL
database server

Primary region

Secondary region

Perform a failover from the primary to the secondary region

Alan Rodrigues

Azure SQL database - Auto-failover



Azure SQL database



Azure SQL database



Azure SQL database server



Azure SQL database server

Primary region

Secondary region

Here you can replicate and failover a set of databases from a primary region to a secondary region.

This feature is built on top of geo-replication.

Here the databases can failover as a group

You can initiate the failover or have a policy that can carry out the failover process.

Your application also does not need to change the connection string when it comes to the failover.



Azure SQL database



Azure SQL database server

Primary region



Azure SQL database



Azure SQL database server

Secondary region

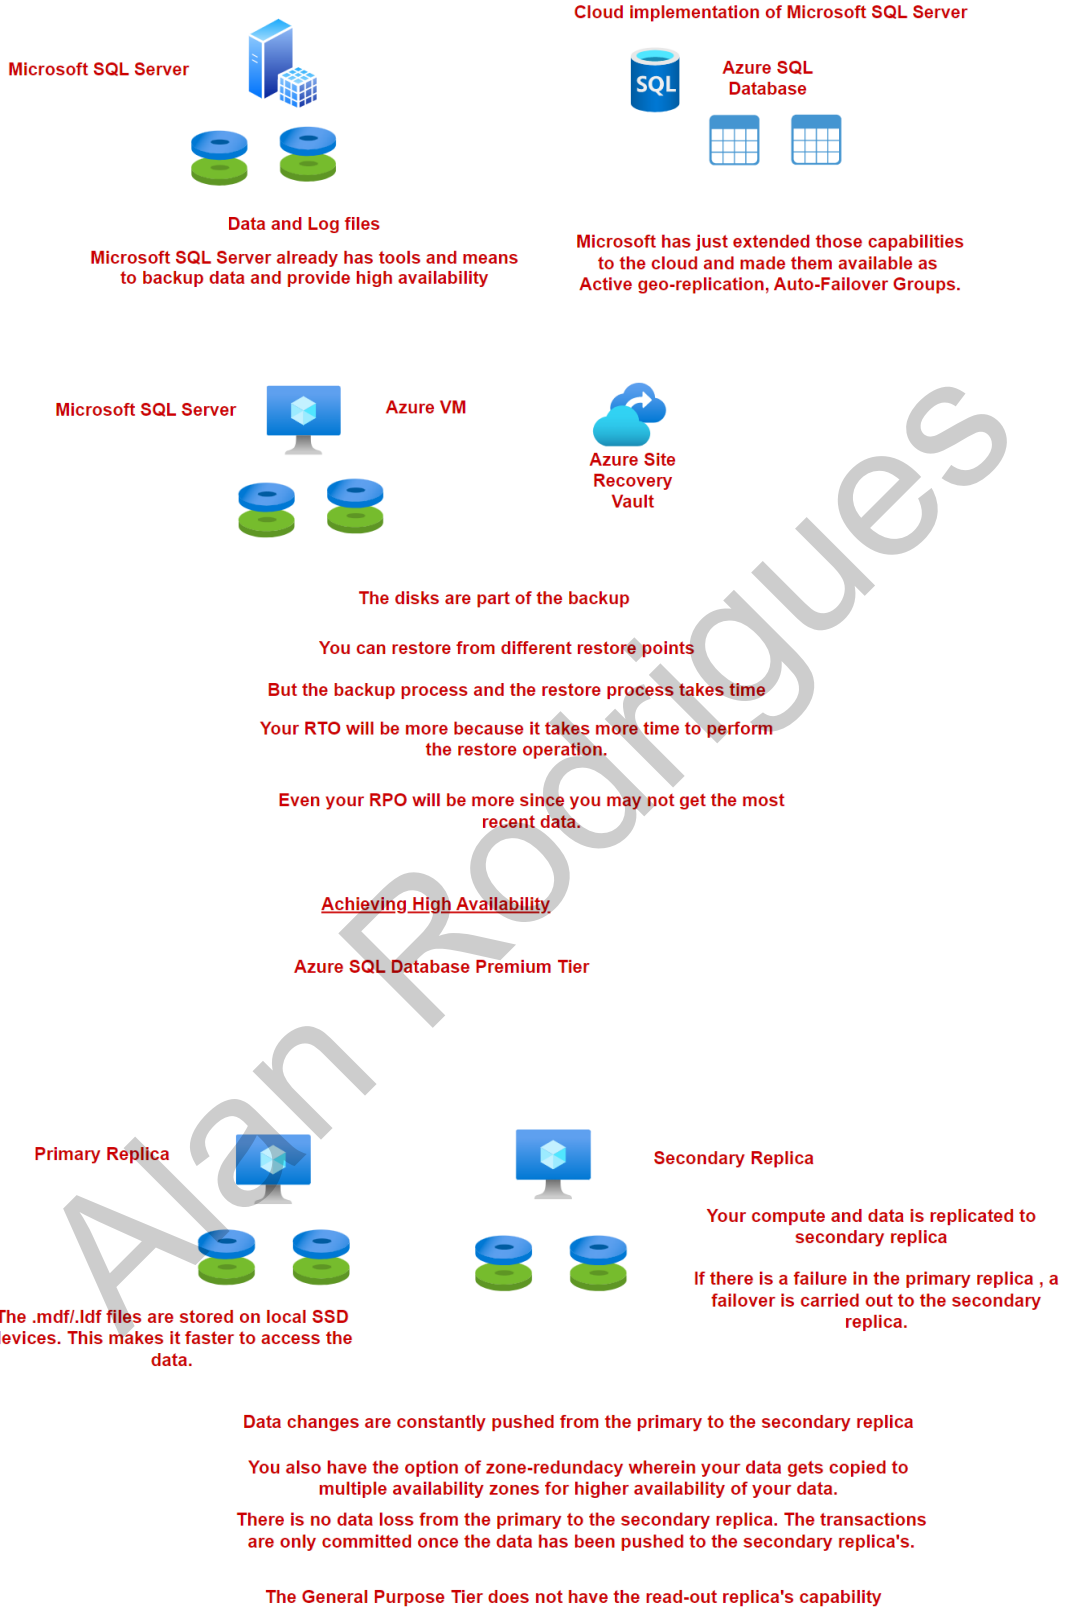
Listener endpoints

app

After the failover the DNS record is automatically updated to send requests to the endpoints in the secondary region.

Alan Rodrigues

Review on Azure SQL database availability



Service and compute tier

Select from the available tiers based on the needs of your workload. The vCore model provides a wide range of configuration controls and offers Hyperscale and Serverless to automatically scale your database based on your workload needs. Alternately, the DTU model provides set price/performance packages to choose from for easy configuration. [Learn more](#)


Service tier
[Compare service tiers](#)

Compute tier
 Provisioned - Compute resources are pre-allocated. Billed per hour based on vCores configured.
 Serverless - Compute resources are auto-scaled. Billed per second based on vCores used.

Compute Hardware

Select the hardware configuration based on your workload requirements. Availability of compute optimized, memory optimized, and confidential computing hardware depends on the region, service tier, and compute tier.

Hardware Configuration **Standard-series (Gen5)**
up to 128 vCores, up to 625 GB memory
[Change configuration](#)



Cost summary

General Purpose (GP, Gen5.2)	
Cost per vCore (in USD)	184.09
vCores selected	x 2
Cost per GB (in USD)	0.12
Max storage selected (in GB)	x 41.6
ESTIMATED COST / MONTH	372.97 USD

Save money

Already have a SQL Server License? Save with a license you already own with Azure Hybrid Benefit. Actual savings may vary based on region and performance tier. [Learn more](#)

Yes No

vCores [Compare vCore options](#)

2

Data max size (GB)

32

9.6 GB LOG SPACE ALLOCATED

Would you like to make this database zone redundant? Yes No

But you can make the database zone-redundant

For the Business Critical Tier, you have all options but is more expensive

Service and compute tier


Select from the available tiers based on the needs of your workload. The vCore model provides a wide range of configuration controls and offers Hyperscale and Serverless to automatically scale your database based on your workload needs. Alternately, the DTU model provides set price/performance packages to choose from for easy configuration. [Learn more](#)

Service tier
[Compare service tiers](#)

Compute Hardware

Select the hardware configuration based on your workload requirements. Availability of compute optimized, memory optimized, and confidential computing hardware depends on the region, service tier, and compute tier.

Hardware Configuration **Standard-series (Gen5)**
up to 128 vCores, up to 625 GB memory
[Change configuration](#)



Cost summary

Business Critical (BC, Gen5.2)	
Cost per vCore (in USD)	495.99
vCores selected	x 2
Cost per GB (in USD)	0.25
Max storage selected (in GB)	x 15.6
ESTIMATED COST / MONTH	995.88 USD

Save money

Already have a SQL Server License? Save with a license you already own with Azure Hybrid Benefit. Actual savings may vary based on region and performance tier. [Learn more](#)

Yes No

vCores [Compare vCore options](#)

2

Data max size (GB)

12

Save money

Already have a SQL Server License? Save with a license you already own with Azure Hybrid Benefit. Actual savings may vary based on region and performance tier. [Learn more](#)

Yes No

vCores [Compare vCore options](#)

2

Data max size (GB) ⓘ

12

3.6 GB LOG SPACE ALLOCATED

Read scale-out

Enabled Disabled

Would you like to make this database zone redundant? ⓘ

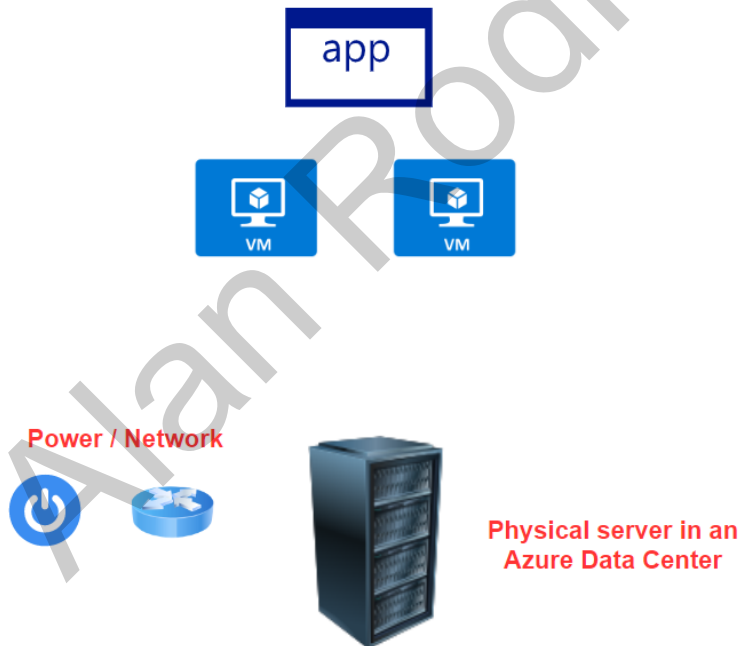
Yes No

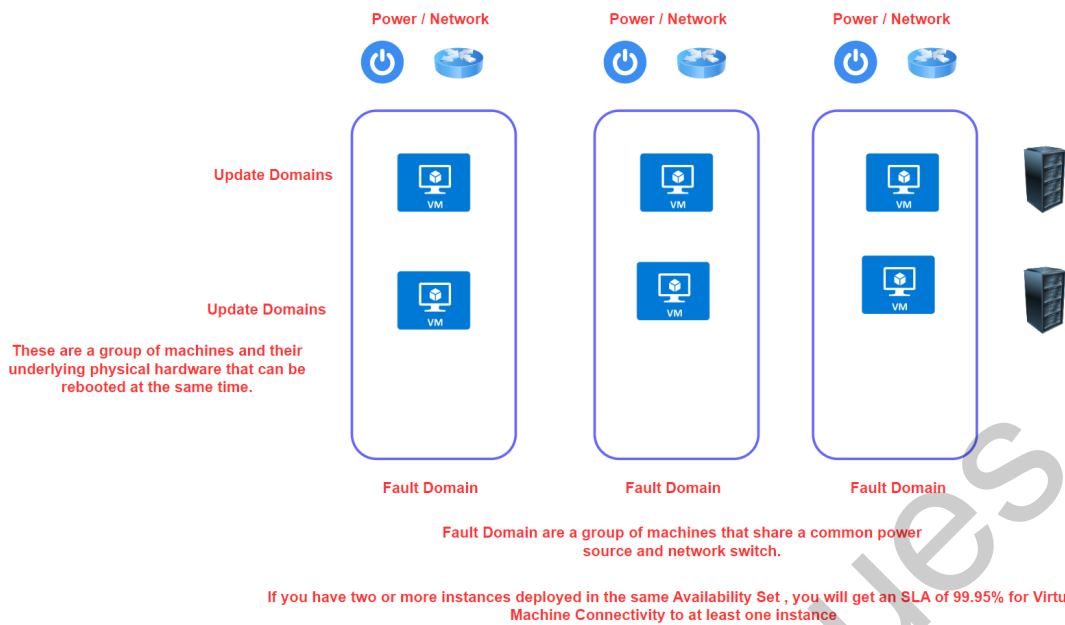
The Serverless compute model is available for the General Purpose and Hyperscale tier

But you cannot pause the compute in the Hyperscale tier

You can have replica's in the Hyperscale tier for higher availability

Availability Sets Review

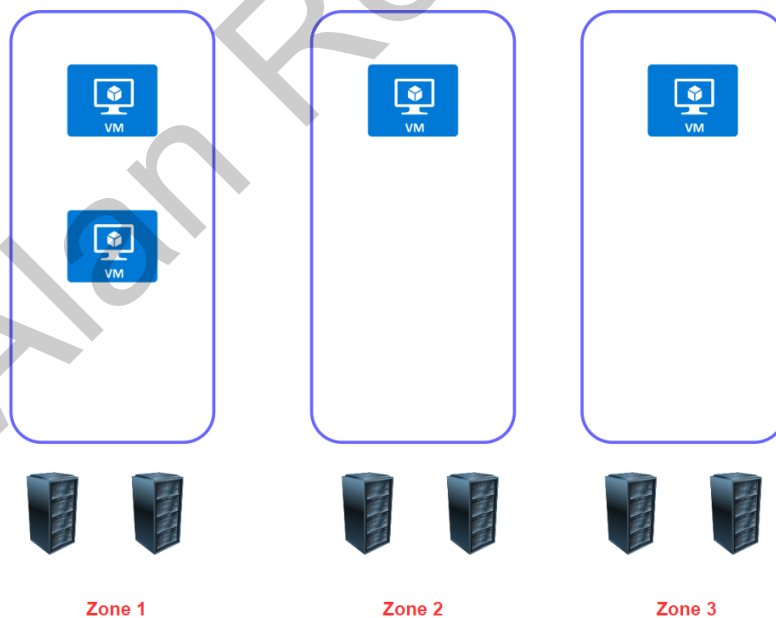




Availability Zones – Review

Availability Zones are unique physical locations that are equipped with independent power, cooling and networking.

There are normally three Availability Zones in a region



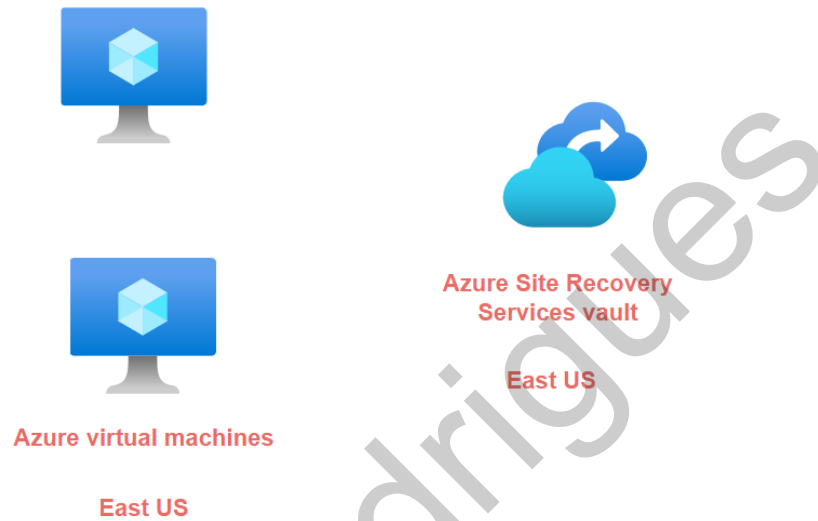
If you have two or more instances deployed in the same Availability Zone , you will get an SLA of 99.99% for Virtual Machine Connectivity to at least one instance

Lab- Azure Backup- Creating the VM

Azure Backup for virtual machines

This provides access to data on the VM if something happens to the original VM

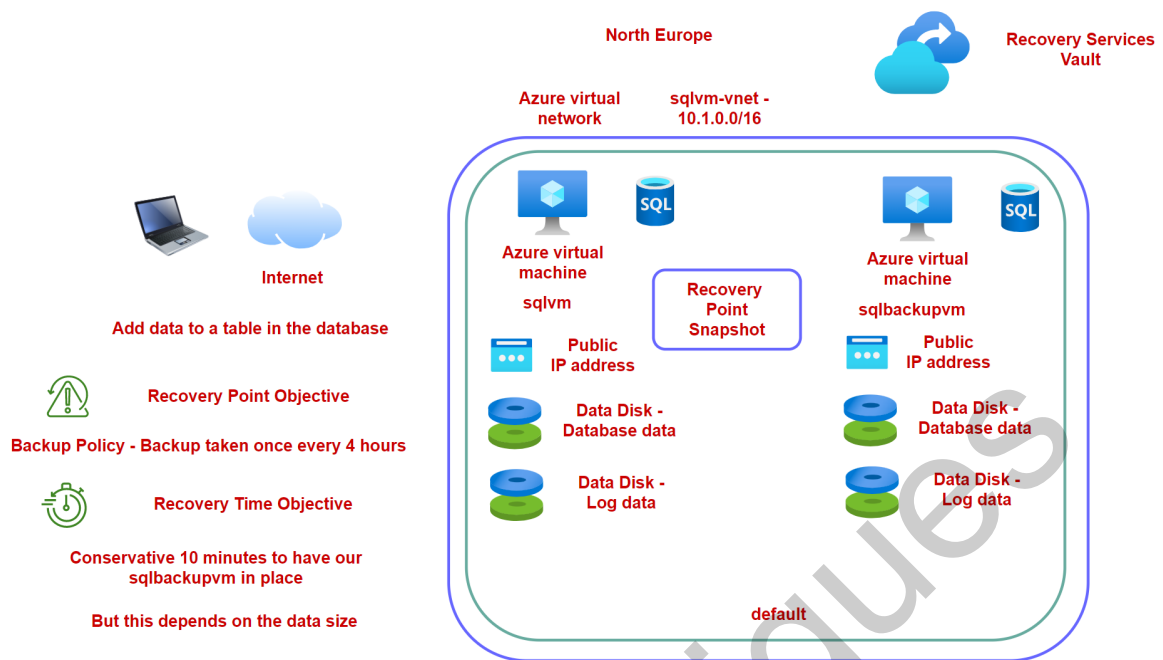
The backup data gets written to a Recovery Services vault



Steps during a backup

1. First an extension is installed on the VM - Supported for both Windows and Linux VM's
2. The backup tool first takes a snapshot of the data and stores it on the local machine
3. The snapshot of data is then copied to the Recovery Service vault

Review on Azure Backup

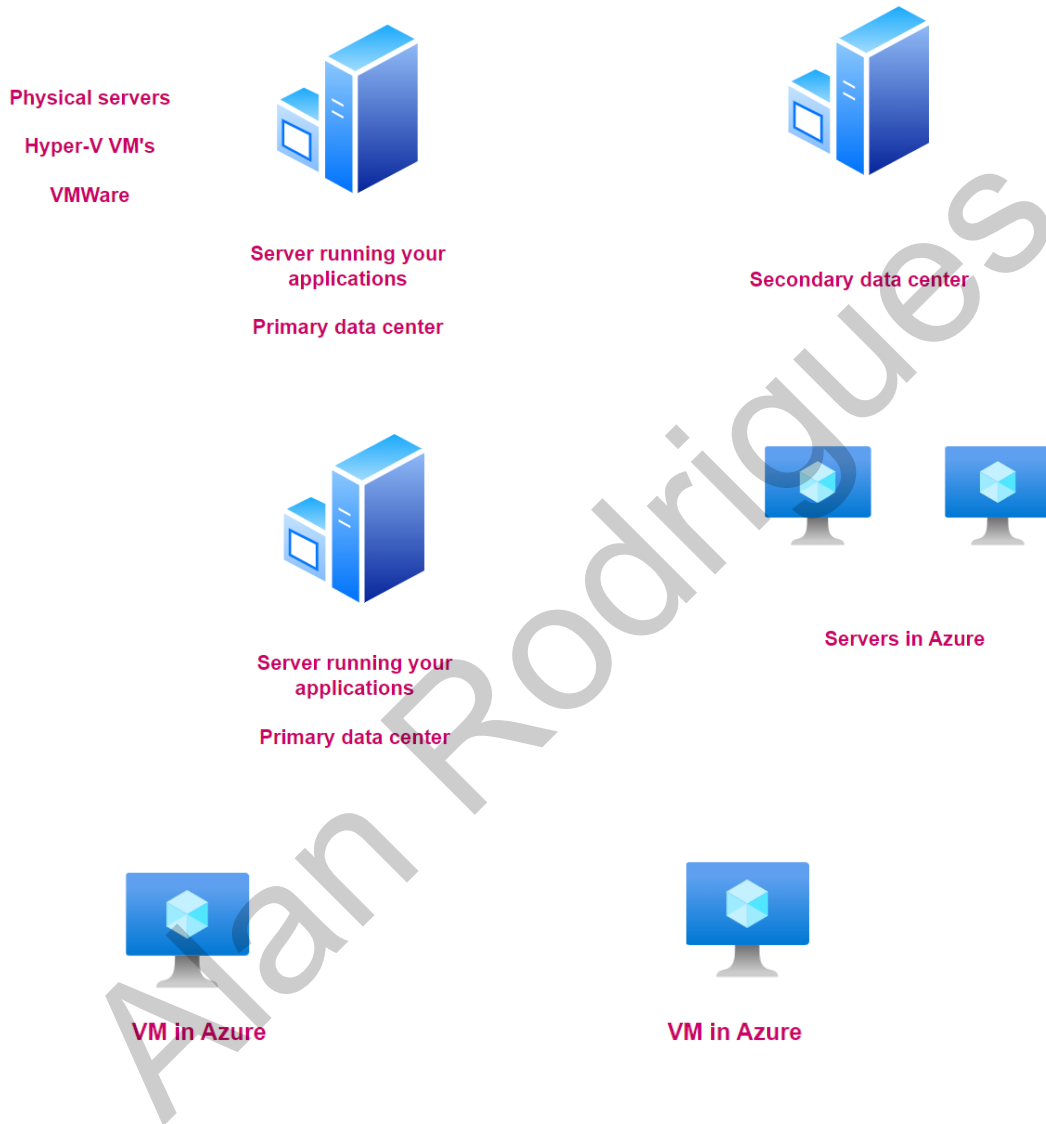


Review- Azure Site Recovery- Virtual Machines

Azure Site Recovery

Used for business continuity and for disaster recovery

Ensures your apps and workloads are running when there are planned or unplanned outages

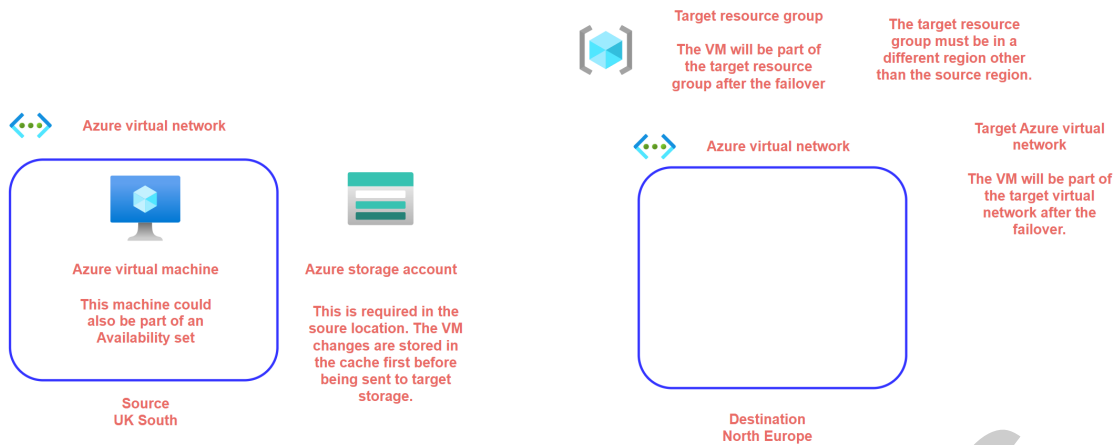


The replication frequency is high , being as low as every 30 seconds for Hyper-V VMs

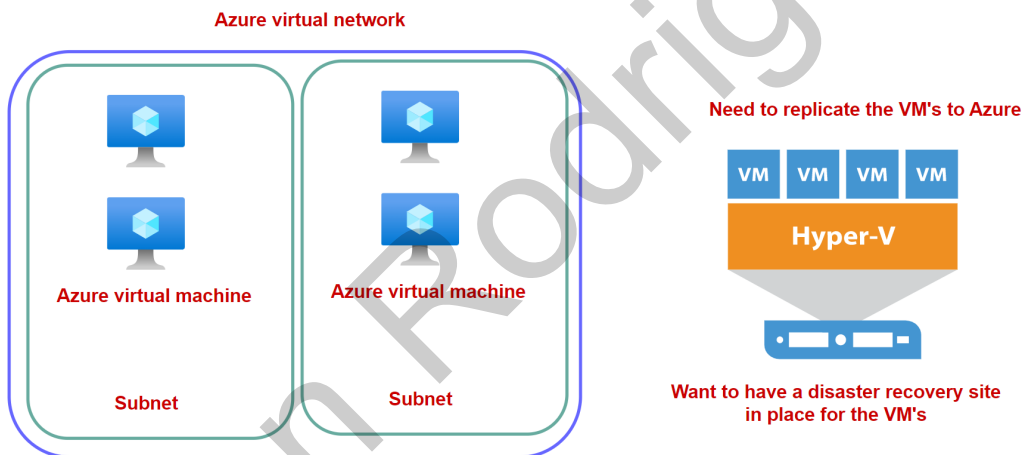
Hence the RPO is low. And because you can switch over quickly, the RTO is also low

You can run planned failovers with zero-data loss

Or unplanned failovers with minimal data loss



Azure Site Recovery – Note



Download and install the Microsoft Azure Site Recovery Provider on the Hyper-V host

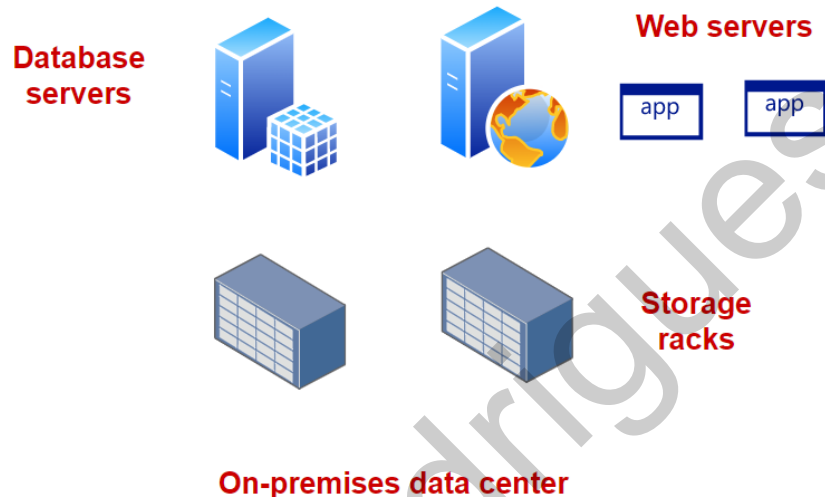
This will enroll the Hyper-V host with the Recovery Services vault

You can then replicate the machines

Design Infrastructure

Migrating your solutions

Migrating solutions



On-premises data center

Migrating web applications

Do you need complete access to the underlying infrastructure?

You want to migrate to Azure Web Apps, is the underlying programming runtime available?

Migrating Microsoft SQL Server

Do you need all of the capabilities of Microsoft SQL Server?

Then you might need to host SQL Server on an Azure VM

Or use Azure SQL Managed Instance if you need managed features of the Azure SQL database service.

Creating new solutions

Should we redesign and convert our solutions to use modern development practices?

Create container-based solutions and use container-based compute solutions?

Can we convert some data stores to NoSQL and use cloud services accordingly?

Transferring data

Do we need to transfer our data to Azure Storage?

Do we have a lot of data to transfer?

What's going to be our strategy to transfer the data?

Azure Virtual Machine



Infrastructure as a service

You get complete control over the underlying VM

You get complete administrative privilege and can install any application on the VM

You have to maintain the underlying compute machine

You can install almost any type of workload

Here scaling needs to be implemented via the use of Virtual Machine scale sets

Azure Web App



Platform as a service

Here the underlying compute VM's are managed by Azure

You can't install anything on the underlying VM's

Support for only web applications. Depends on the underlying supported framework.

You get features such as scaling for your web apps

Azure VM's vs Azure Web Apps

Azure Virtual Machine



Infrastructure as a service

You get complete control over the underlying VM

You get complete administrative privilege and can install any application on the VM

You have to maintain the underlying compute machine

You can install almost any type of workload

Here scaling needs to be implemented via the use of Virtual Machine scale sets

Azure Web App



Platform as a service

Here the underlying compute VM's are managed by Azure

You can't install anything on the underlying VM's

Support for only web applications. Depends on the underlying supported framework.

You get features such as scaling for your web apps

Azure Web Apps- Deployment Slots

Deployment Slots

Staging Environments for App Service Plans



Version 1

Version 2



Production Slot

Staging slot

Standard , Premium and Isolated App Service Plan

Applications in deployment slots have their own host names

1. You have the chance to validate all application changes in the staging deployment slot
2. You can then swap the staging slot with the production slot
3. This helps eliminate the downtime for your application when new changes are deployed
4. You can also easily roll back the changes

Azure Virtual Desktop

Azure Virtual Desktop

This is a desktop and application virtualization service

Setup a multi-session experience for Windows 11 or Windows 10



Azure virtual network



Windows 10



Windows 11

Host pool - This is a collection of Azure virtual machines that need to register onto Azure virtual desktop.

You can create a Personal host pool. This is where each session host is assigned to an individual. This is like providing a dedicated desktop to a user.

You can create a Pool host pool. Here user sessions can be distributed across the session host. This provides a shared experience and can help to save on costs.

You can define an Application group. This is the logical grouping of applications that are installed on the session hosts.

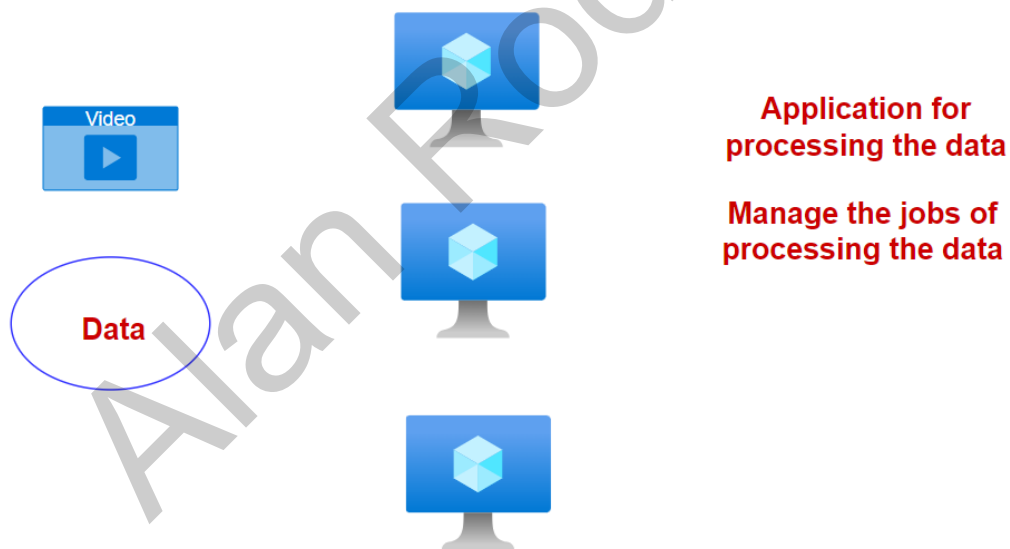
You can also automate certain tasks such as scaling the number of session hosts based on demand.

For this you make use of an Azure Automation Account, a PowerShell runbook and an Azure Logic App.

What is Azure Batch

Azure Batch

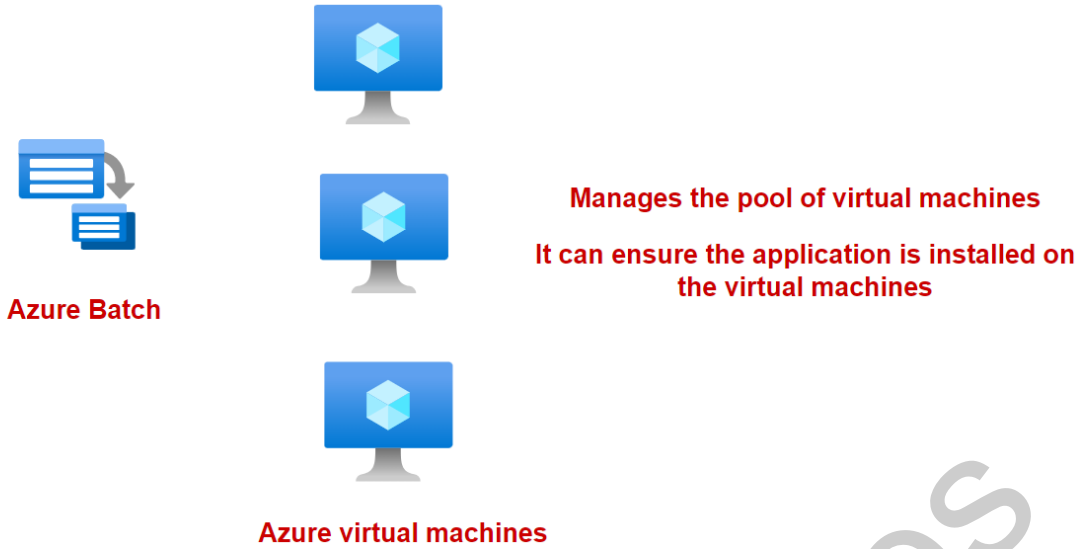
This is used to run large-scale parallel and high performance computing batch jobs



Azure virtual machines



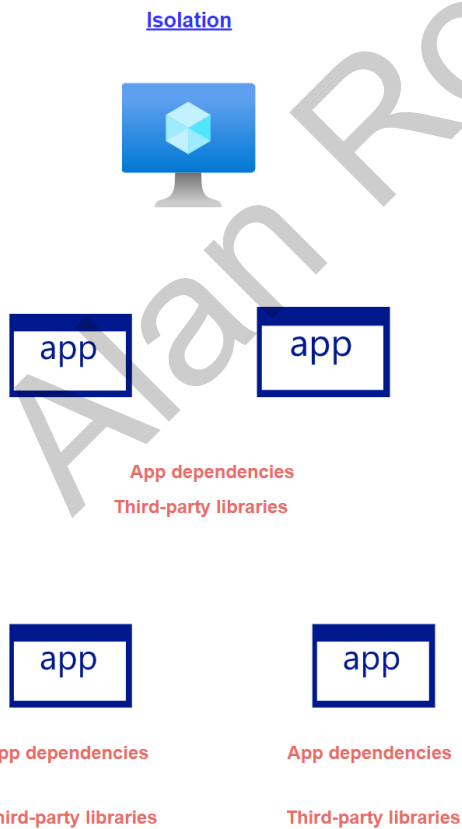
You can manage these batch processing jobs with the help of the Azure Batch service



You can use Azure Storage accounts for the input files, storage of the application and the output files.

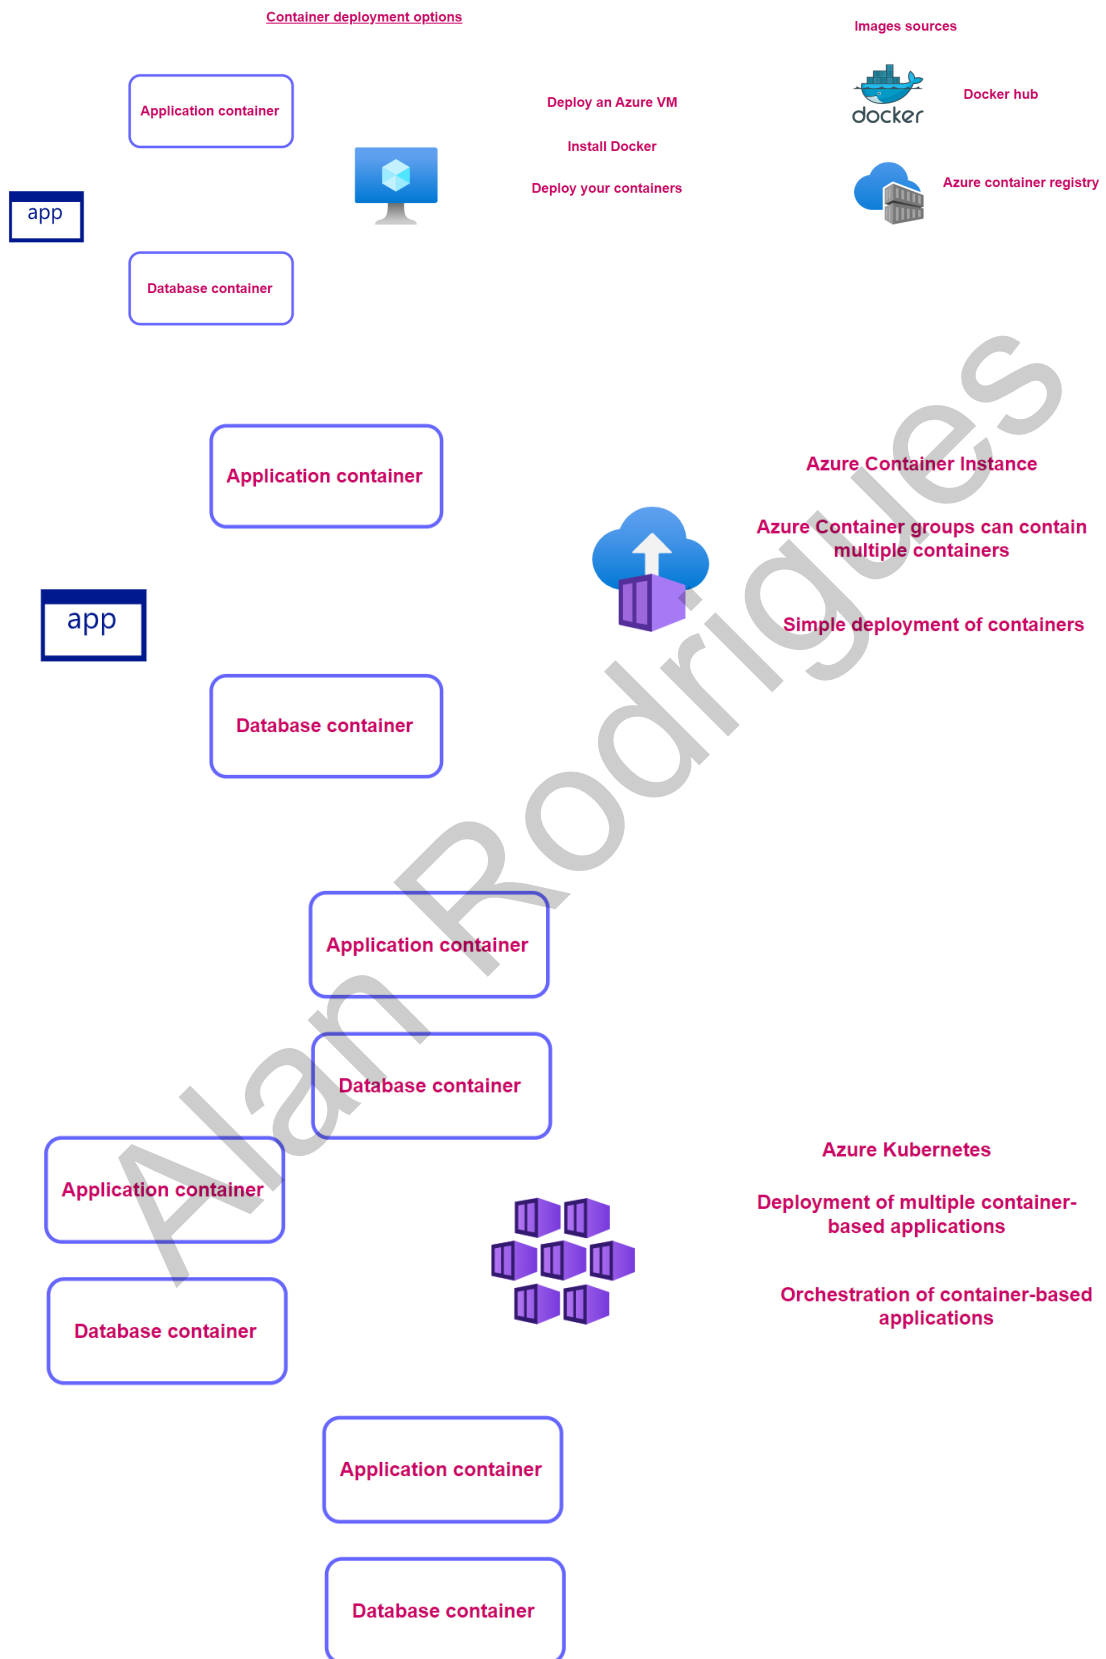
Azure Batch manages the tasks and the jobs that need to run on the virtual machines.

The need for containers



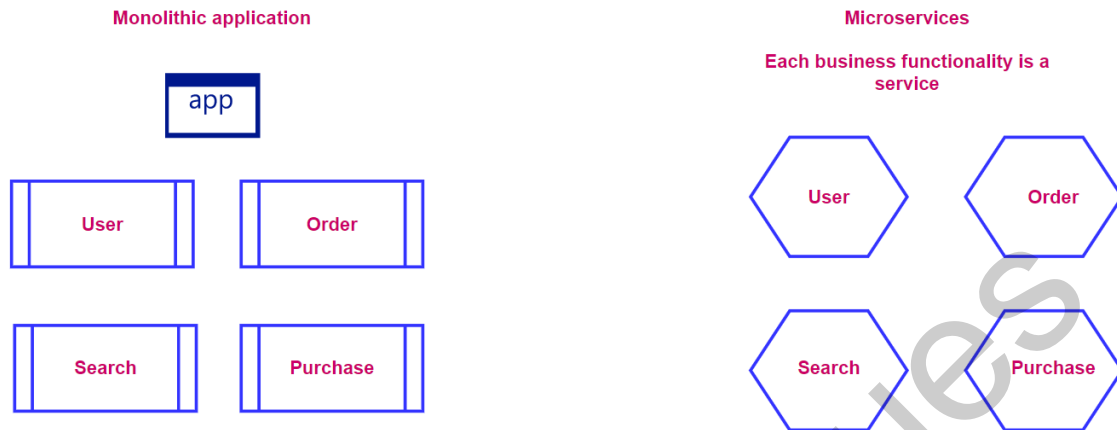
Containers helps to package the application along with libraries , frameworks and dependencies that are required.

The different container deployment options



Microservices Architecture

Microservices Architecture



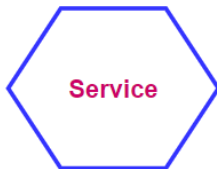
Each service is self-contained

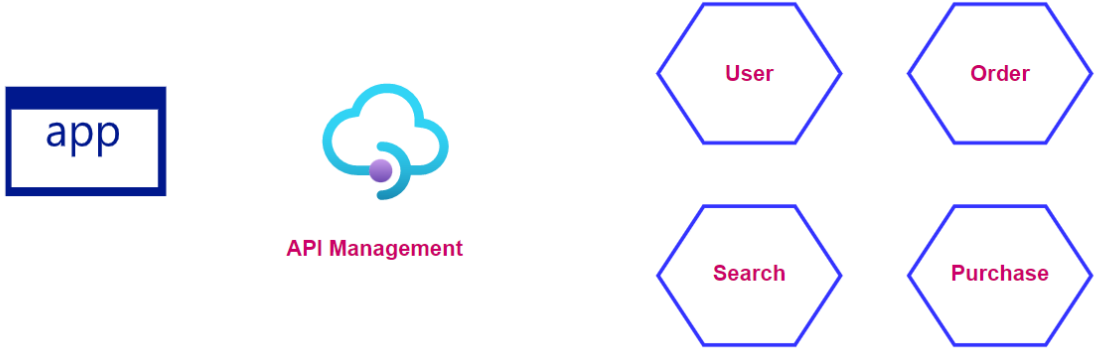
Each service implements a specific business function

Each service has a separate codebase that is managed by a small development team

It should be easy to update the service independently without breaking the entire application that consumes the service

Each service normally persists its own data





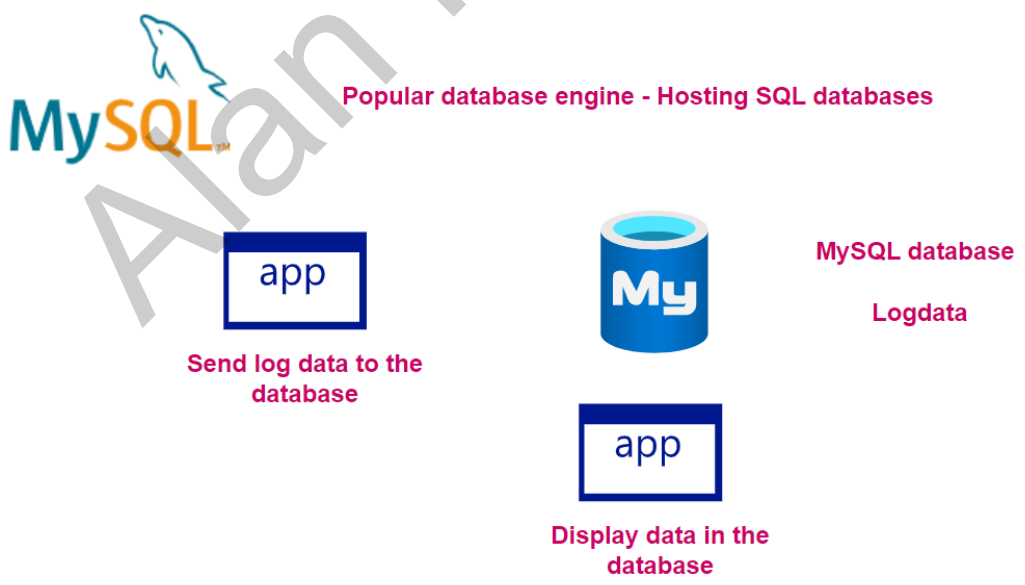
Advantages

- Agile
- Smaller code base
- Mix and match technologies
- Fault isolation

Disadvantages

- Complexity
- Testing
- Governance
- Data Integrity

Working with containers- What we want to achieve



Option 1



Display data in the database



Install MySQL

Create a database

Create the table to store application data

Option 2



Display data in the database



Use Azure database for MySQL

Platform as a service option

What are we going to start with



Send log data to the database



Use Azure database for MySQL

Platform as a service option



Display data in the database

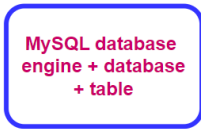
What we want to accomplish



Send log data to the database



Use Azure database for MySQL
Platform as a service option



Run the database in a container



Display data in the database



Run our application that displays data in a container



Linux VM
Docker host

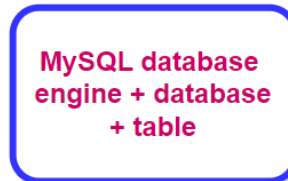
We need a custom environment to build our images



Azure Container registry



**Deploy the custom MySQL
container to an Azure
container instance**



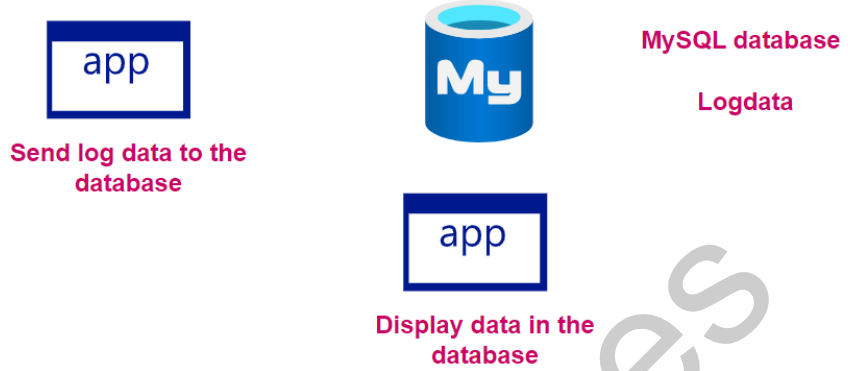
**Deploy both containers to
Azure Kubernetes**

Lab- Let's build a custom MySQL image

Alan Rodrigues

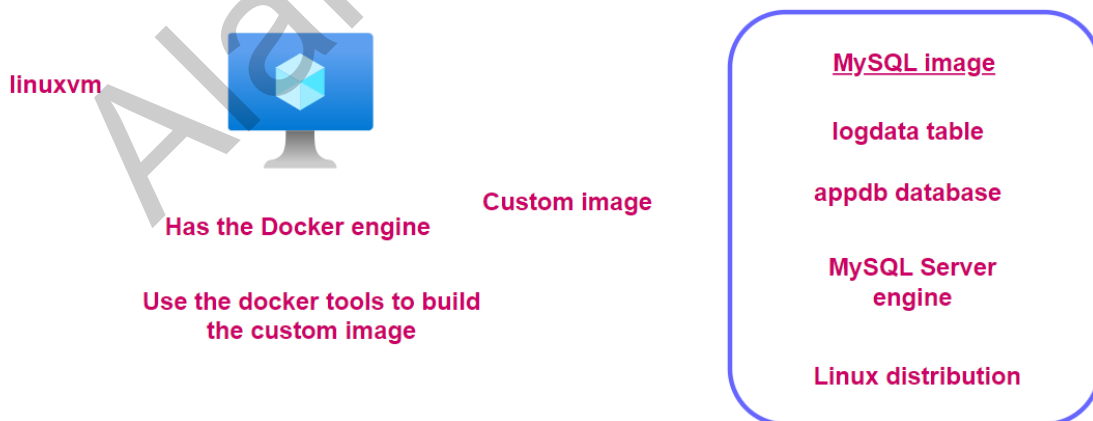
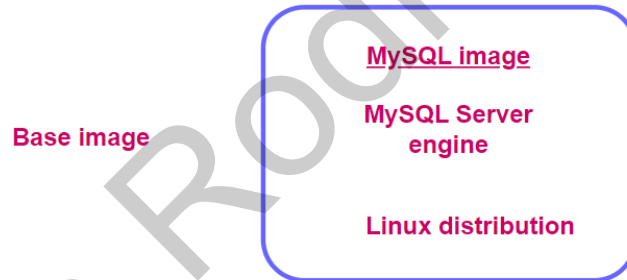


Popular database engine - Hosting SQL databases

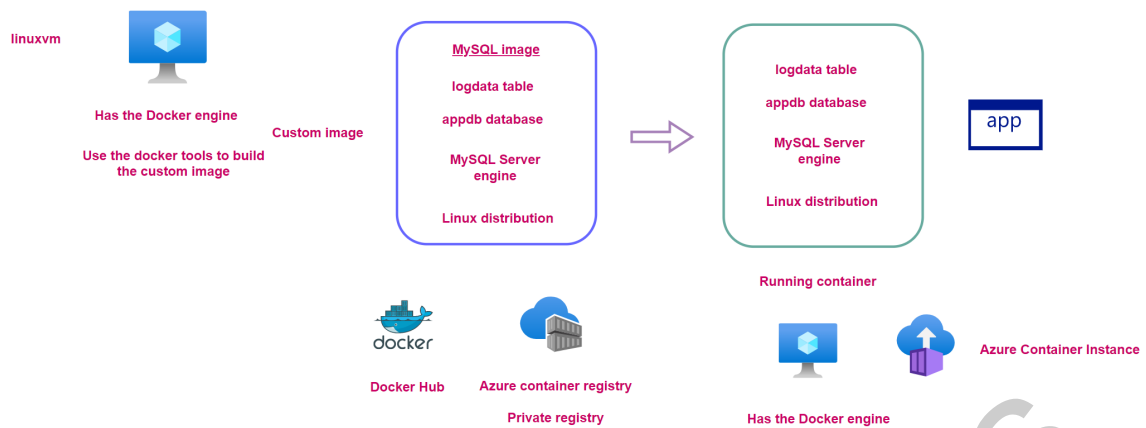


We want our database to run in a container

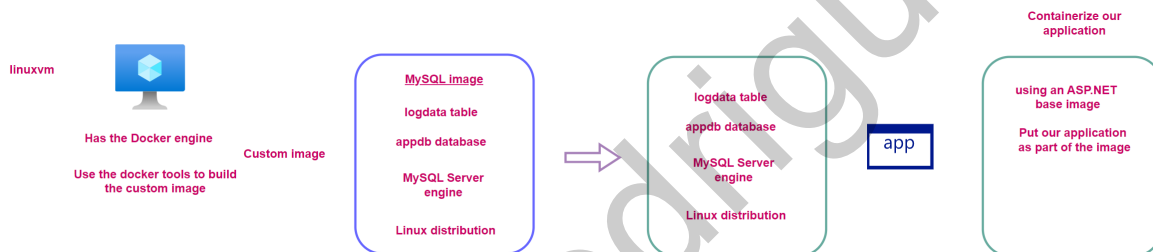
We are using the mysql image that is available in DockerHub or the docker.io image registry



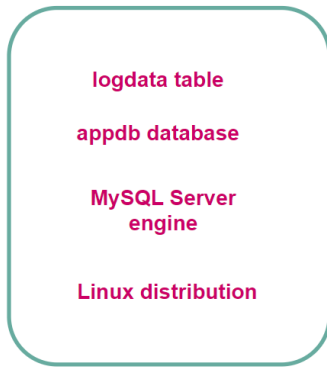
Lab- Deploy the MySQL image to Azure container instance



Lab- Let's build our application image



Azure Container Instances- Other features



Azure Container Instance

Container Restart Policy


Always - Containers in the container group are always restarted
OnFailure - Restarted only when the process executed in the container fails

Containers can also persist data to an Azure file share. If the container restarts the data on the container is lost.

Using Kubernetes

Alan Rodrigues

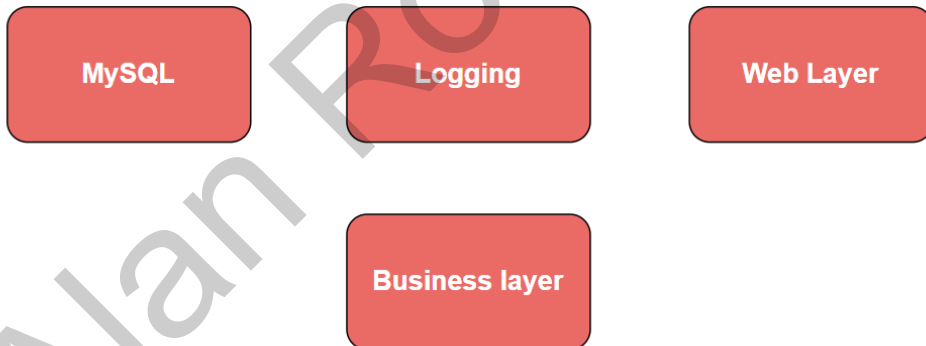
Kubernetes



Managing containers at scale

Azure Kubernetes - Managed service for Kubernetes on Azure

Kubernetes is used to orchestrate your containers for hosting your applications



Azure File Sync

Azure Virtual Machine



Azure File shares

You can create multiple file shares for users

Virtual Machines can also connect to file shares

Azure File Sync



Azure File shares



Azure file sync agent



File server

Headquarters



File server

Branch office



Azure file sync agent

Your on-premises servers can hold all files in the file share

Your on-premises servers can hold a cache of the frequently accessed files in the file share

Cloud tiering



Backup of the file shares

Azure File shares



File server

Headquarters



File server

Branch office

Also if any server goes down, you can just replace the file server

You just have to install the Azure File Sync agent

Review- Virtual Network Peering



staging-network

Virtual network

10.0.0.0/16



test-network

Virtual network

10.1.0.0/16



stagingvm

SubnetA

10.0.0.0/24



testvm

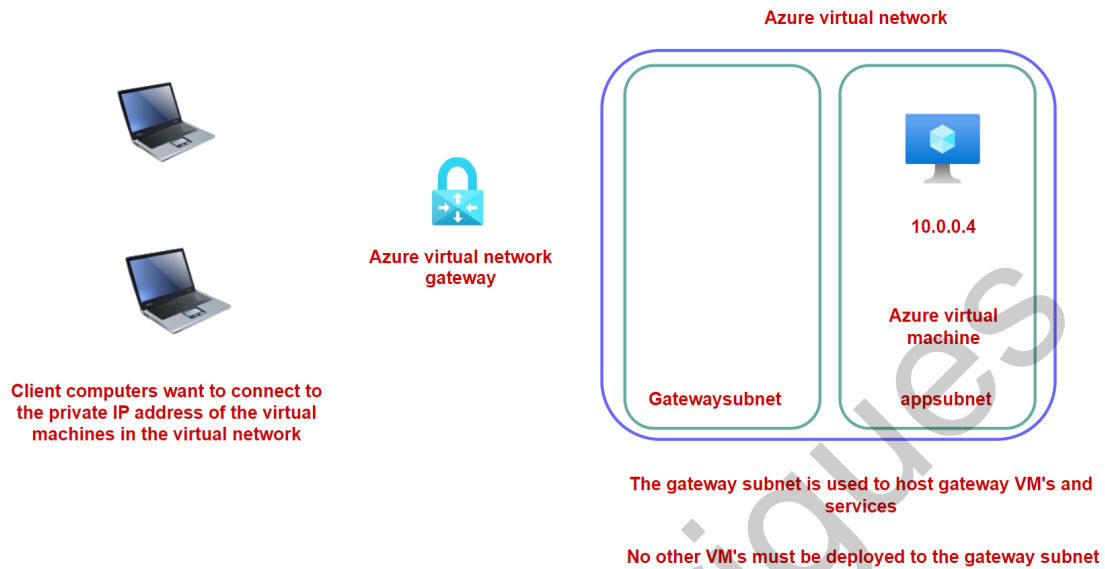
SubnetB

10.0.1.0/24

Review- Connectivity to on-premises networks

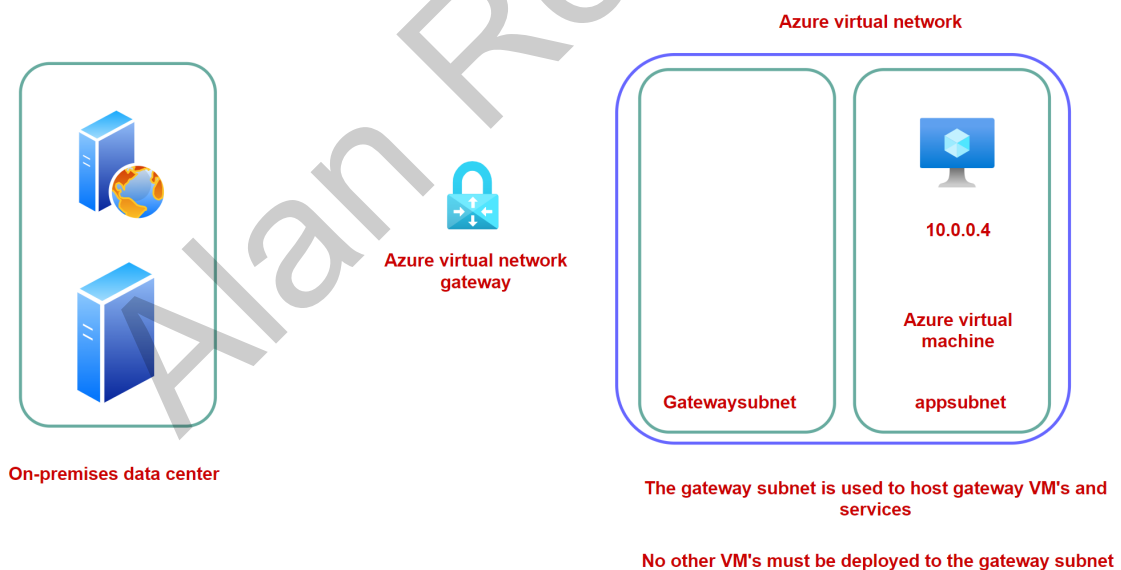
Point-to-Site VPN

This helps you create a secure connection between client computers and your virtual network

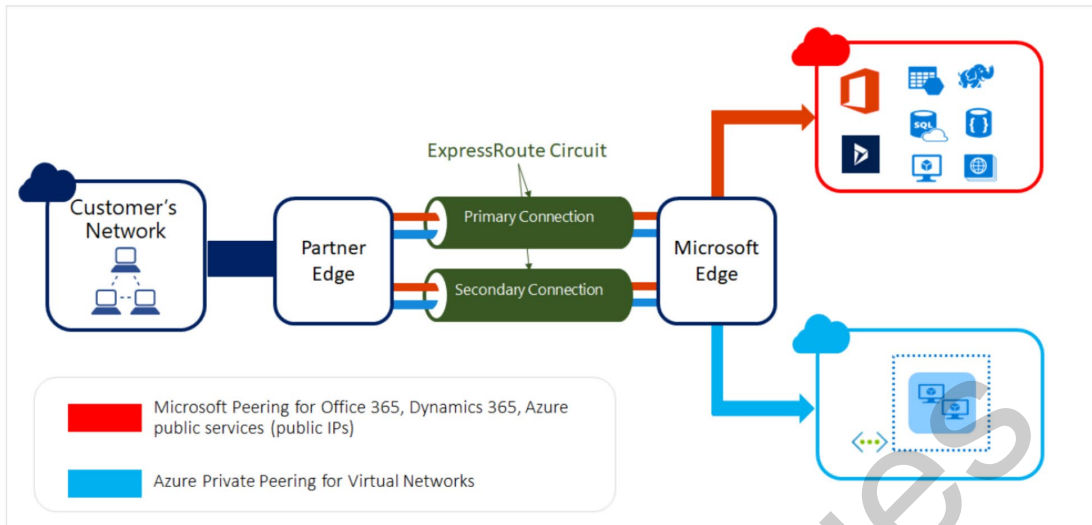


Site-to-Site VPN

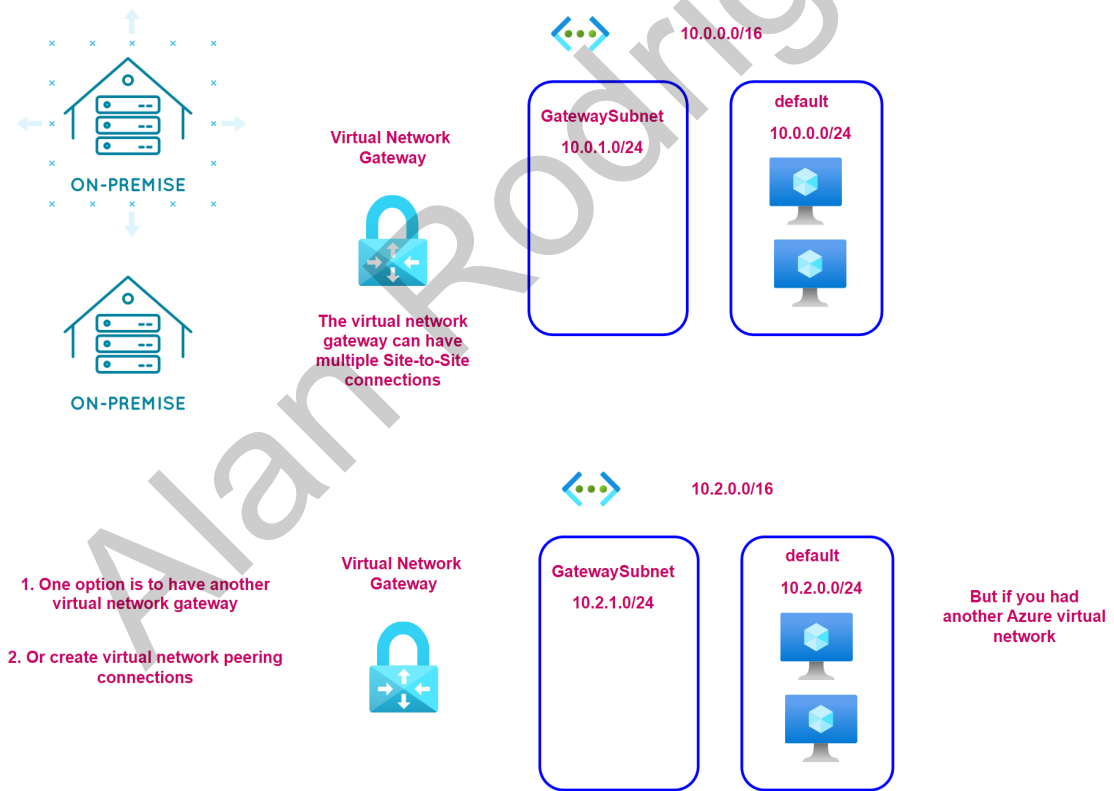
This helps you create a secure connection between your on-premises network and your virtual network

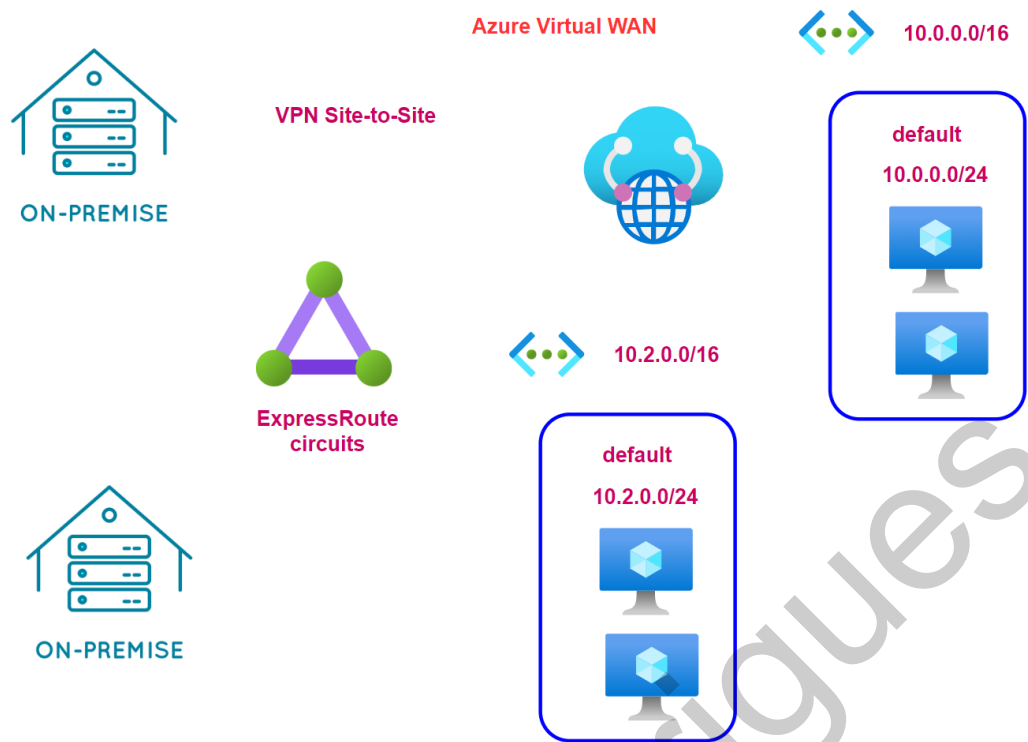


Azure ExpressRoute

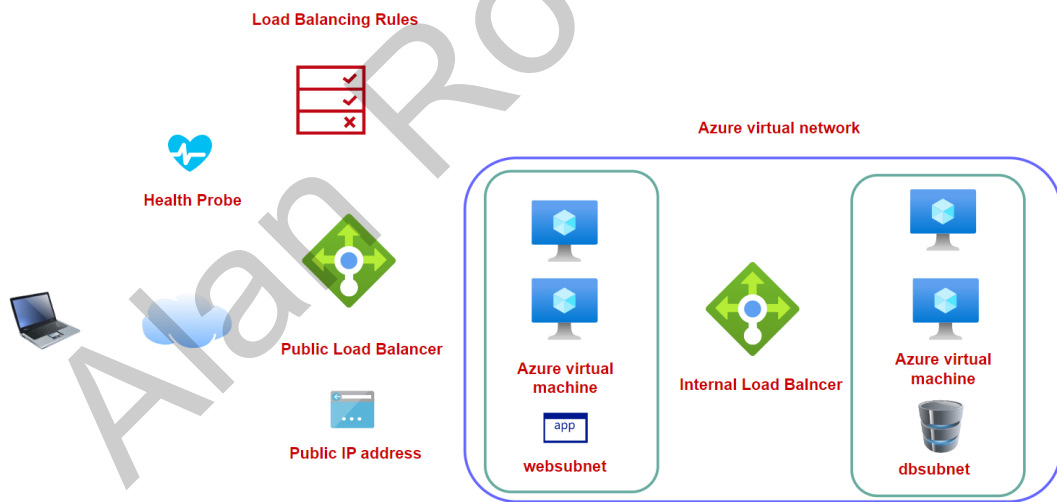


Azure VirtualWAN Review





Azure Load Balancer – Review



Basic Load Balancer

Free

The machines in the backend pool need to be part of an availability set or scale set

Health probes - TCP, HTTP

No support for Availability zones

No SLA

Standard Load Balancer

Charge per hour

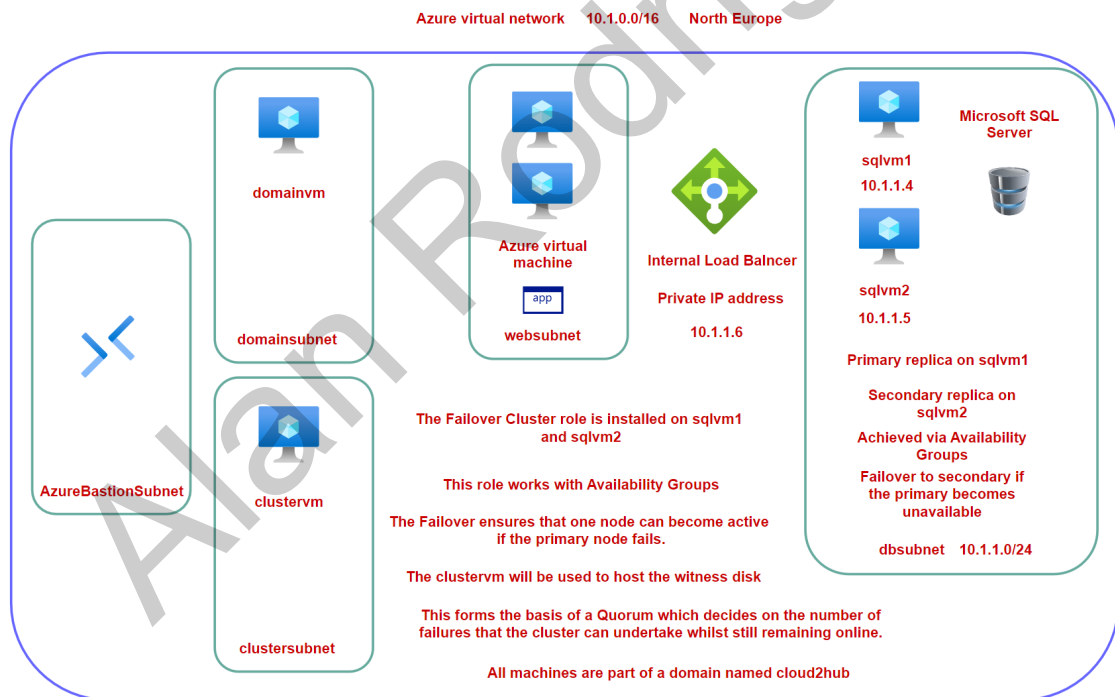
Here the machines can also be independent machines that are part of a virtual network

Health probes - TCP, HTTP, HTTPS

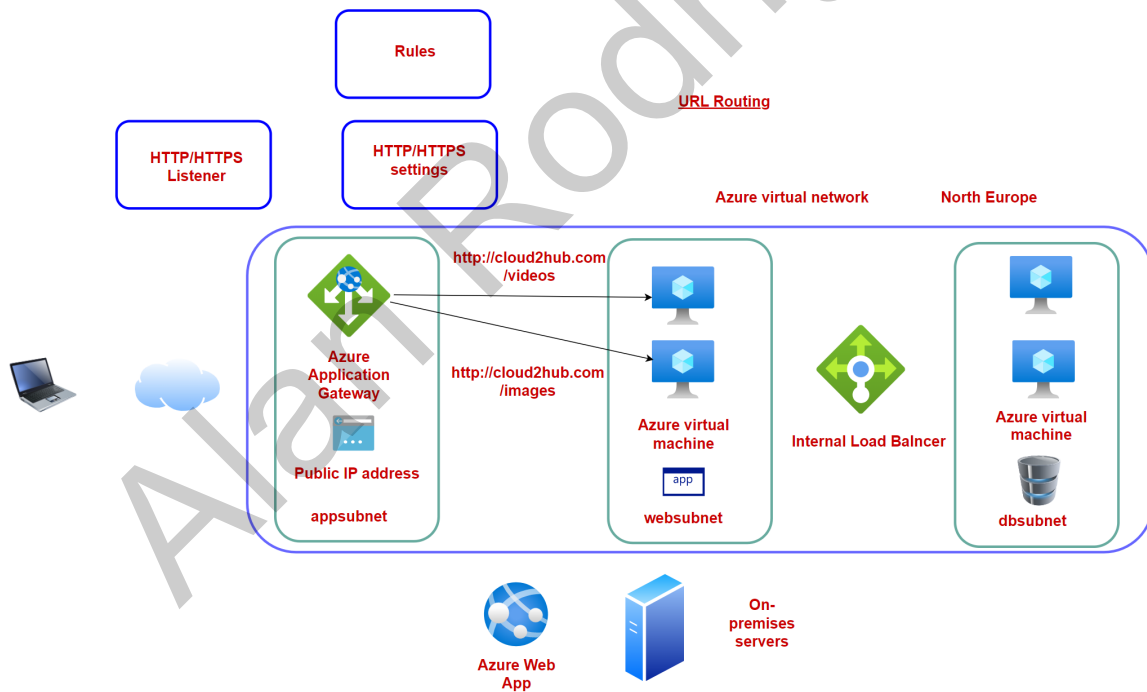
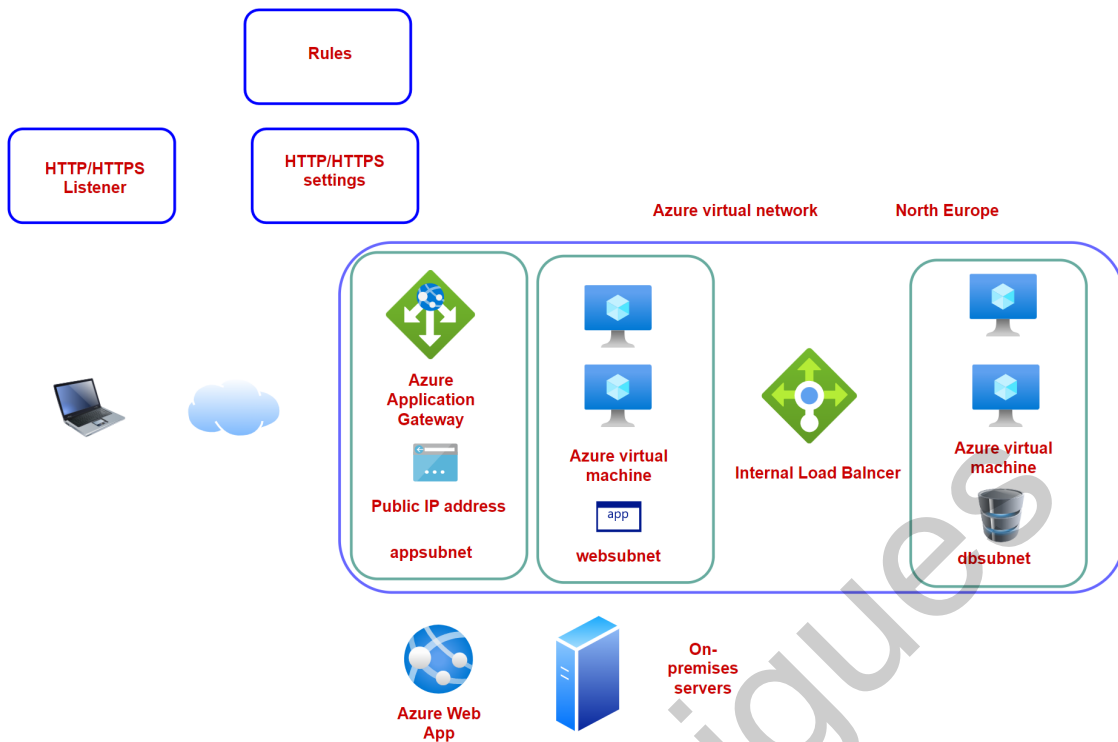
Support for Availability zones

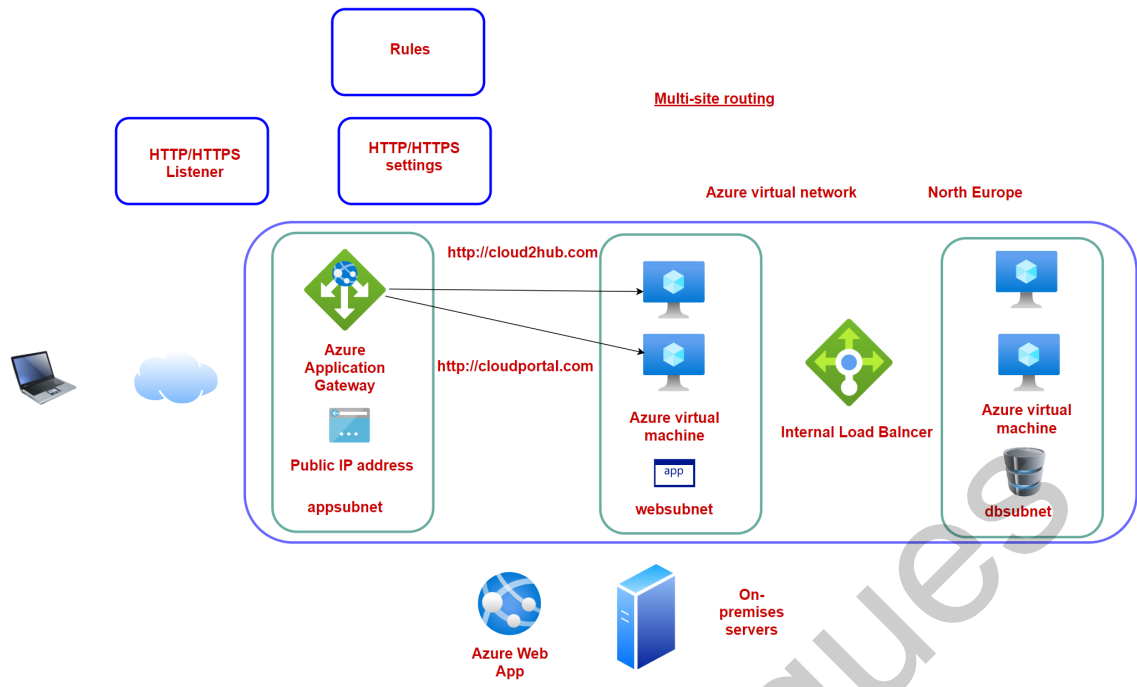
SLA of 99.99%

Azure Internal Load Balancer demo – Design

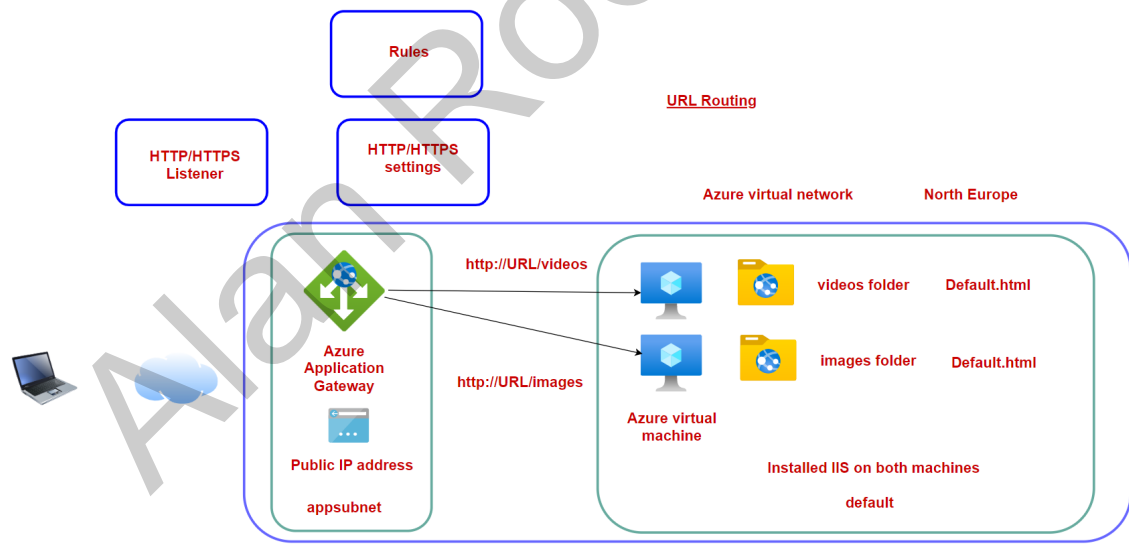


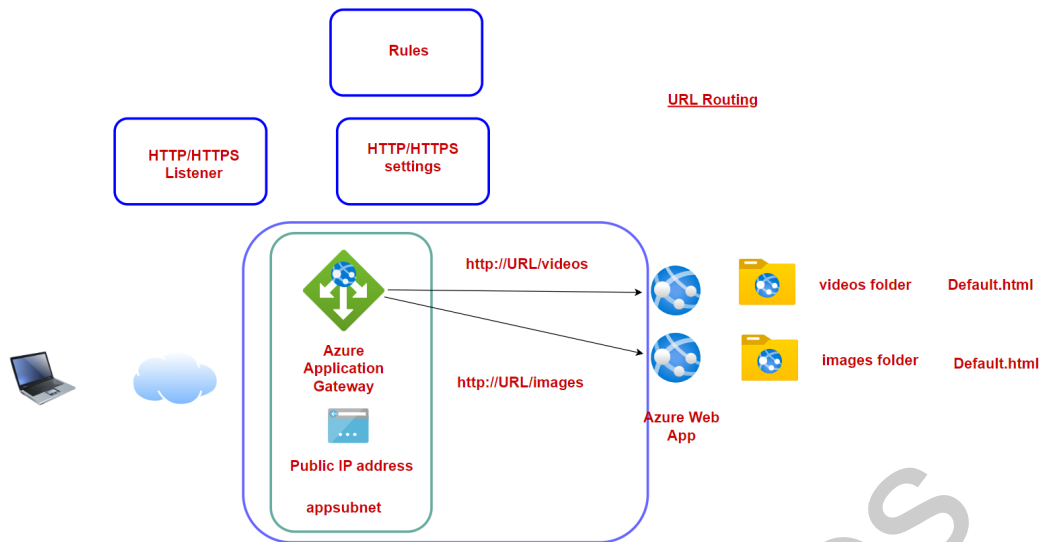
Review of the Azure Application Gateway





Azure Application Gateway- Implementation Review





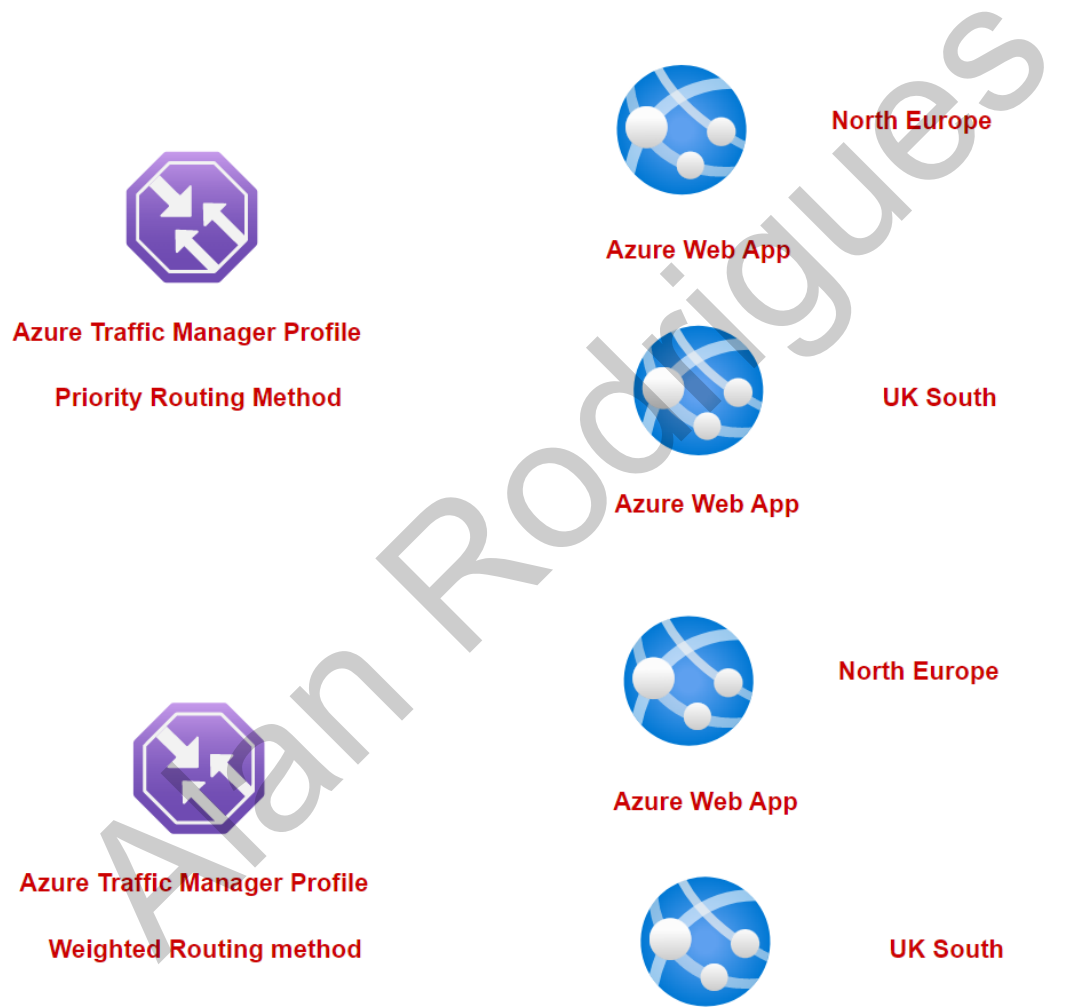
Alan Rodrigues

Azure Traffic Manager

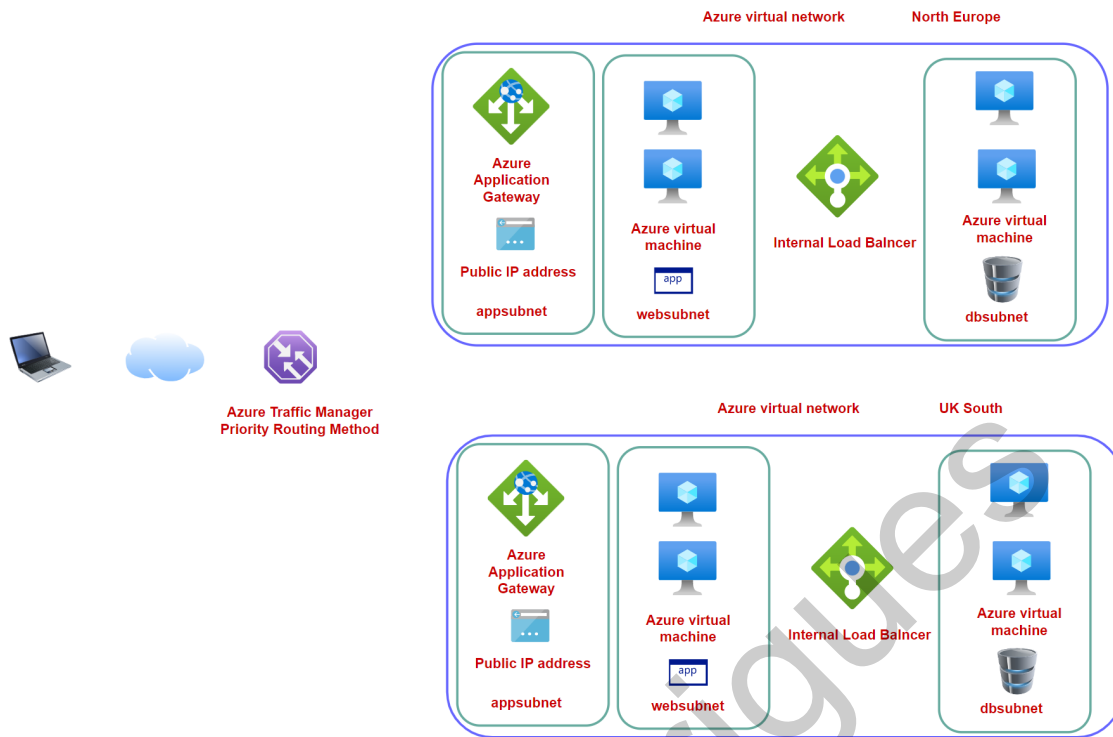
This is a DNS-based traffic load balancer.

You can distribute traffic to public facing applications across different Azure regions.

You can direct traffic based on different routing methods.



Architecture with Traffic Manager

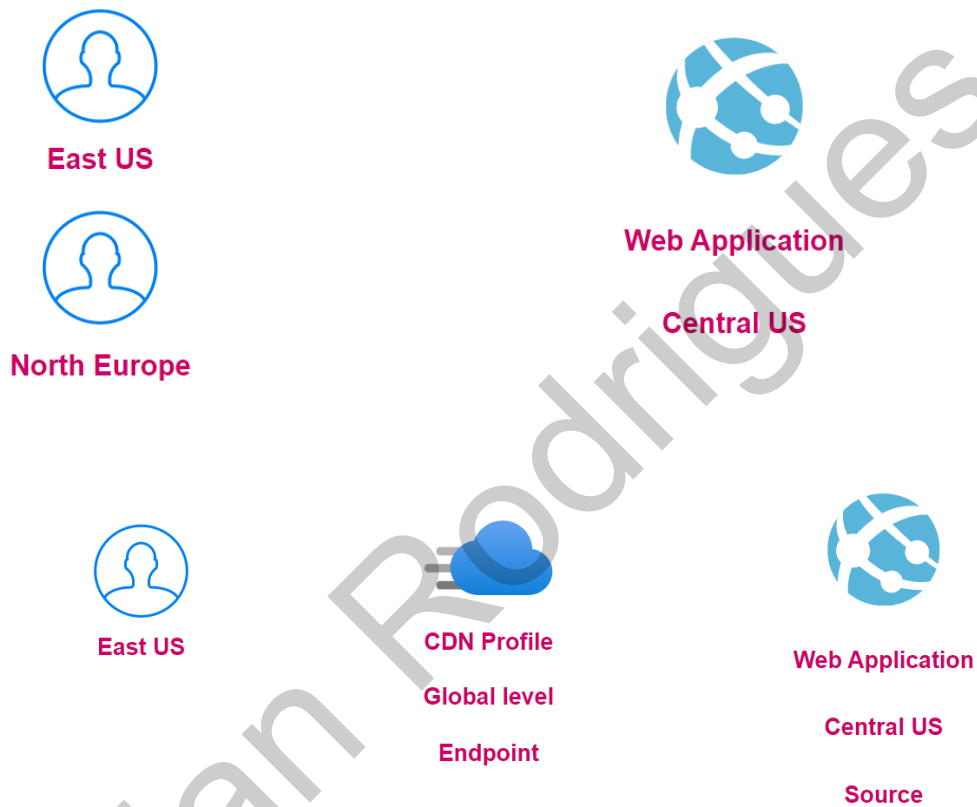


Azure Content Delivery Networking service

Azure CDN

Content Delivery Network

Helps to deliver content to users across the globe by placing content on physical nodes placed across the world



1. The user in the East US location makes a request to the CDN endpoint
2. The CDN checks whether the Point of presence location closest to the user has the requested file.
3. If not a request is made to the source to get the required file.
4. A server in the Point of presence location will then cache the required file.
5. The server will also send the file to the user.
6. Subsequent users from the same location will now be served the file from the server in the point of presence location.

1. Data Cache



Web Application



Azure Cache for Redis



Azure SQL database



Faster memory

Top 10 courses for the day

The application would first calculate the top 10 courses based on the data in the database

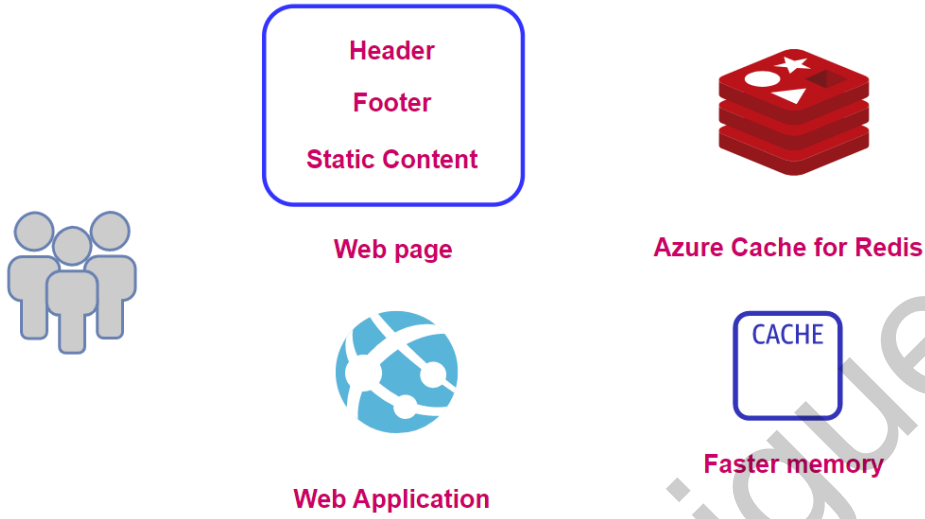
Then the application would store the top 10 courses along with any supporting data to Azure Cache for Redis

The application would then fetch this data for users from Azure Cache for Redis

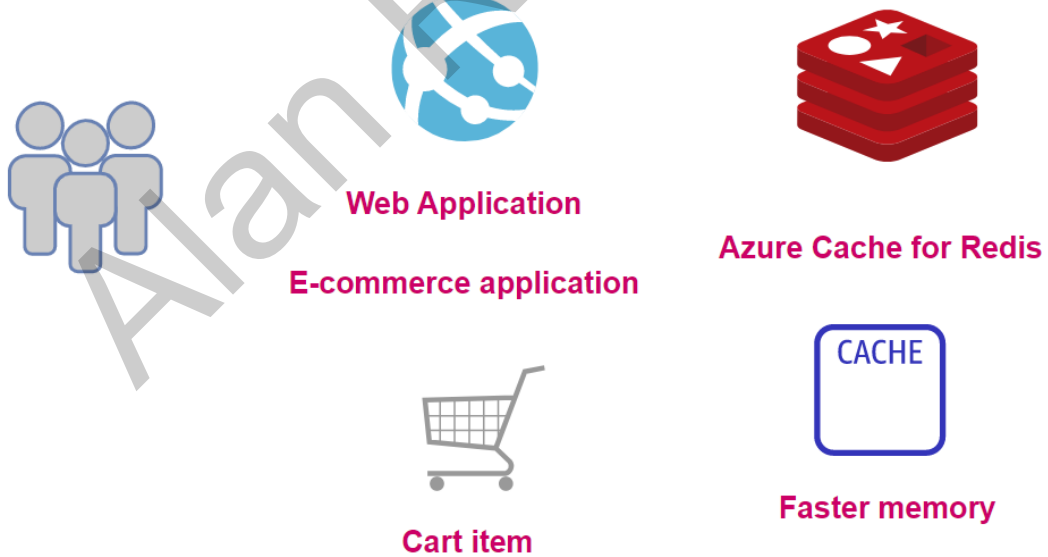
Alan Rodrigues

The application would then update the data in Azure Cache for Redis on a daily basis

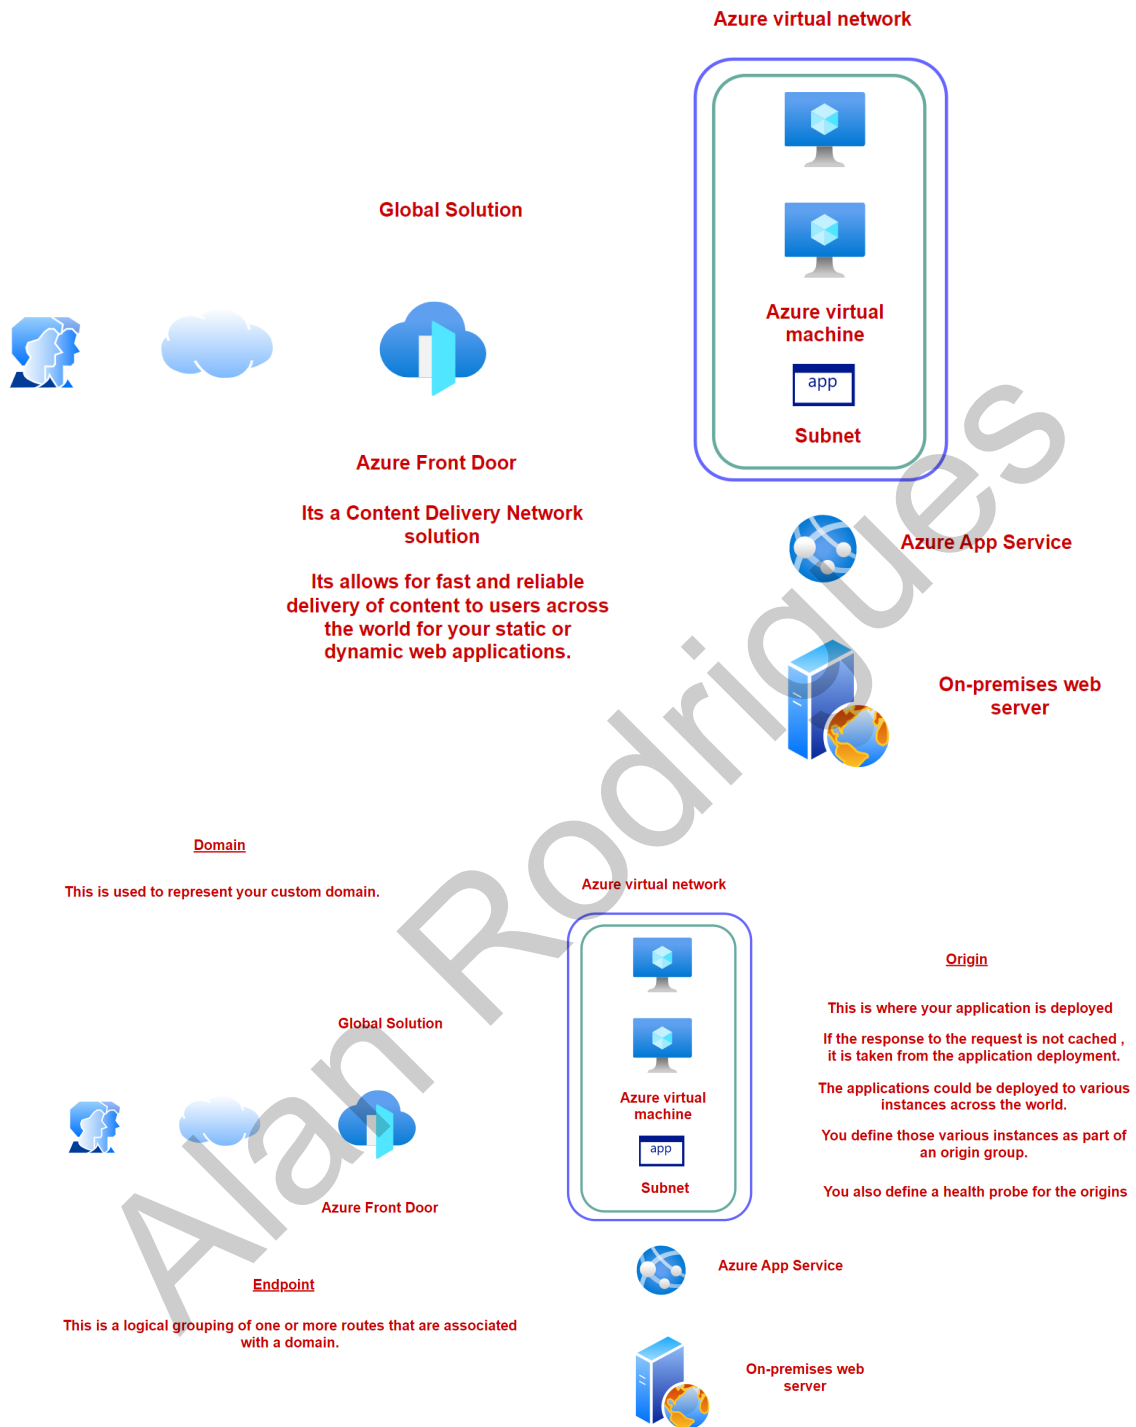
2. Content Cache



3. Session store



Azure Front Door Service – Overview



What are Azure Event Hubs

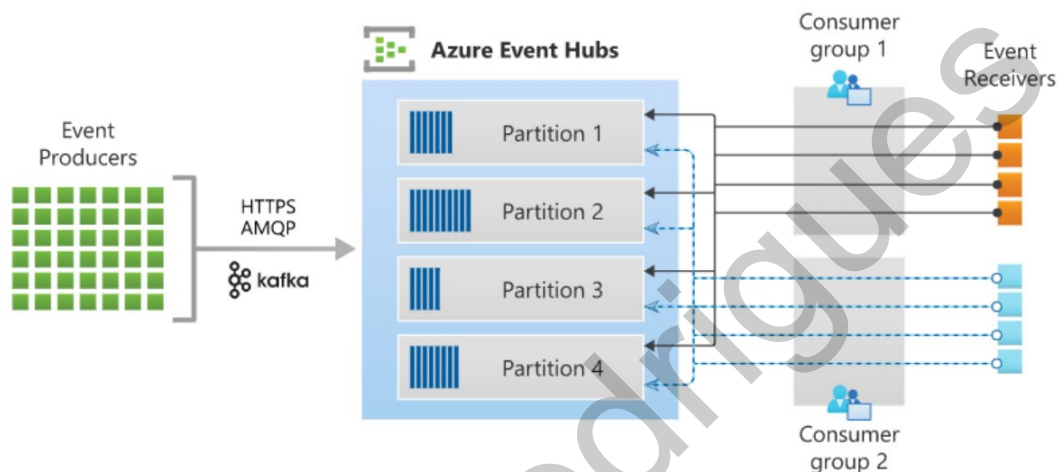
What are Azure Event Hubs

This is a big data streaming platform

This service can receive and process millions of events per second

You can stream log data , telemetry data, any sort of events to Azure Event Hubs

Event Hubs Architecture



<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features>

The different components

Event producers - This is an entity that sends data to an event hub. The events can be published using the protocols - HTTPS, AMQP, Apache Kafka

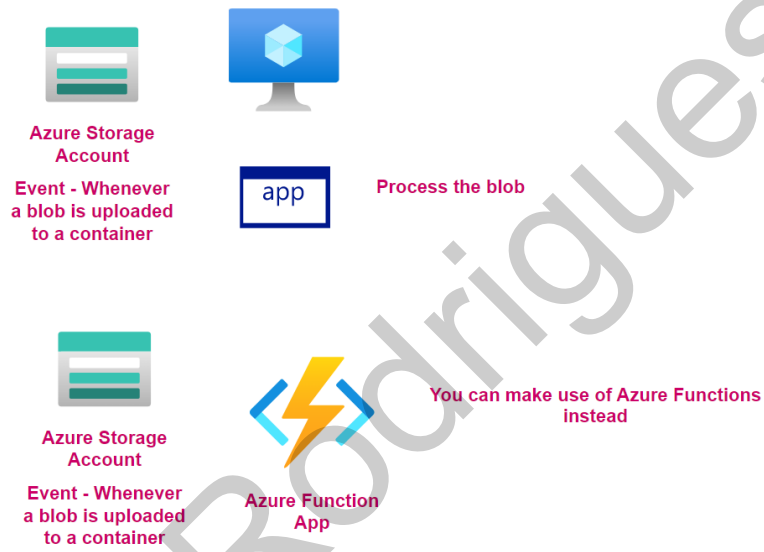
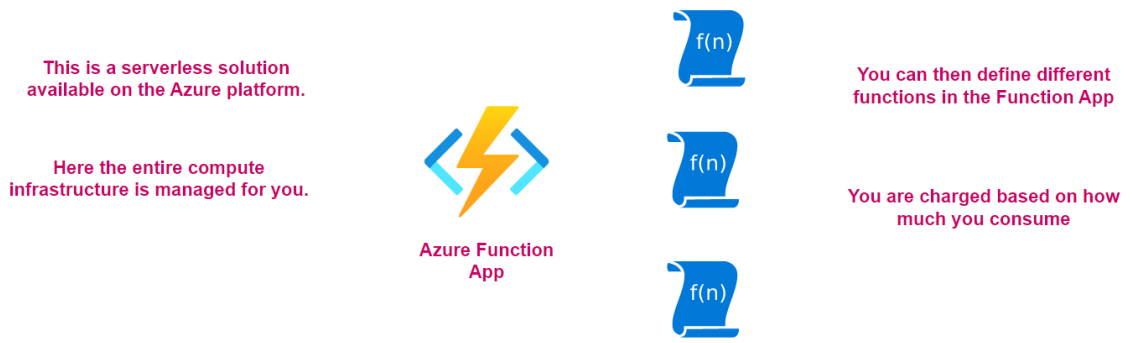
Partitions - The data is split across partitions. This allows for better throughput of your data onto Azure Event Hubs

Consumer groups - This is a view (state, position or offset) of an entire event hub

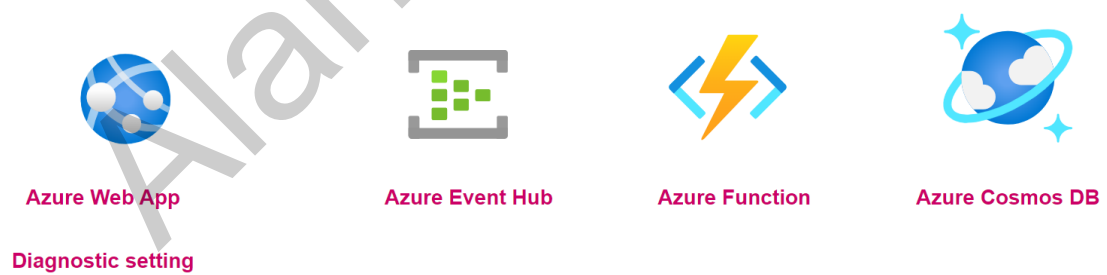
Throughput - This controls the throughput capacity of Event Hubs

Event Receivers - This is an entity that reads event data

Azure Functions



Lab- Architecture scenario- Azure Functions



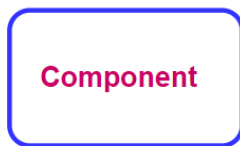
Lab- Azure Service Bus



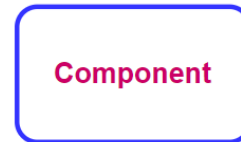
Messaging service

Azure Storage queues

Azure Service Bus



System



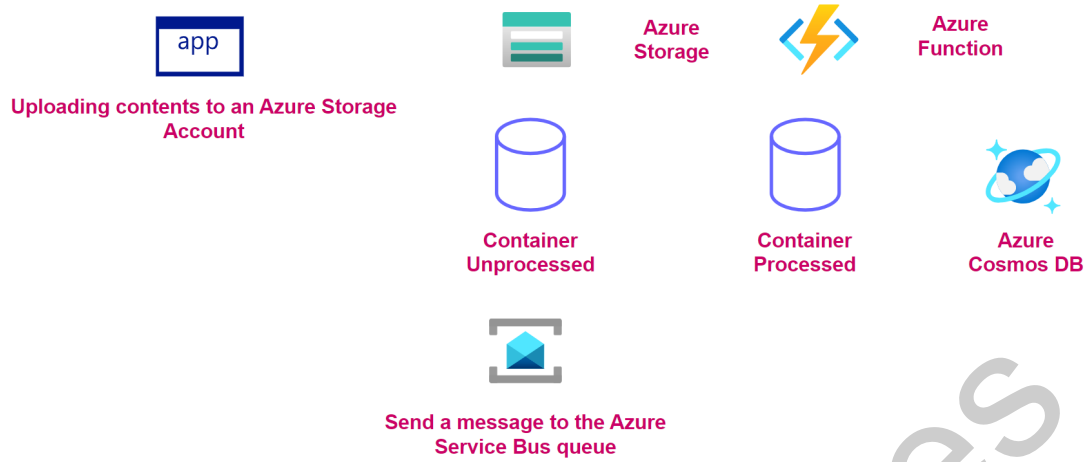
Upload Image



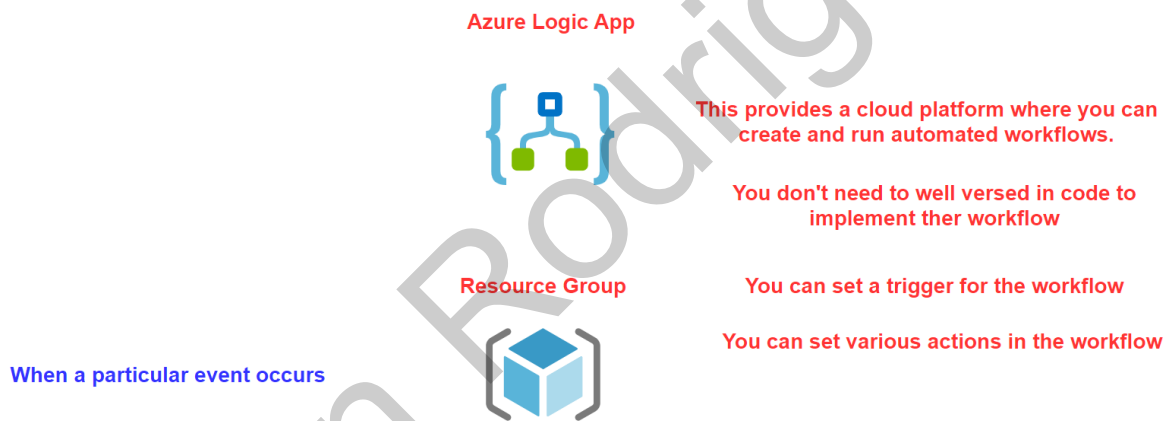
Process Image



Azure Service Bus- Architecture Scenario- Sending data

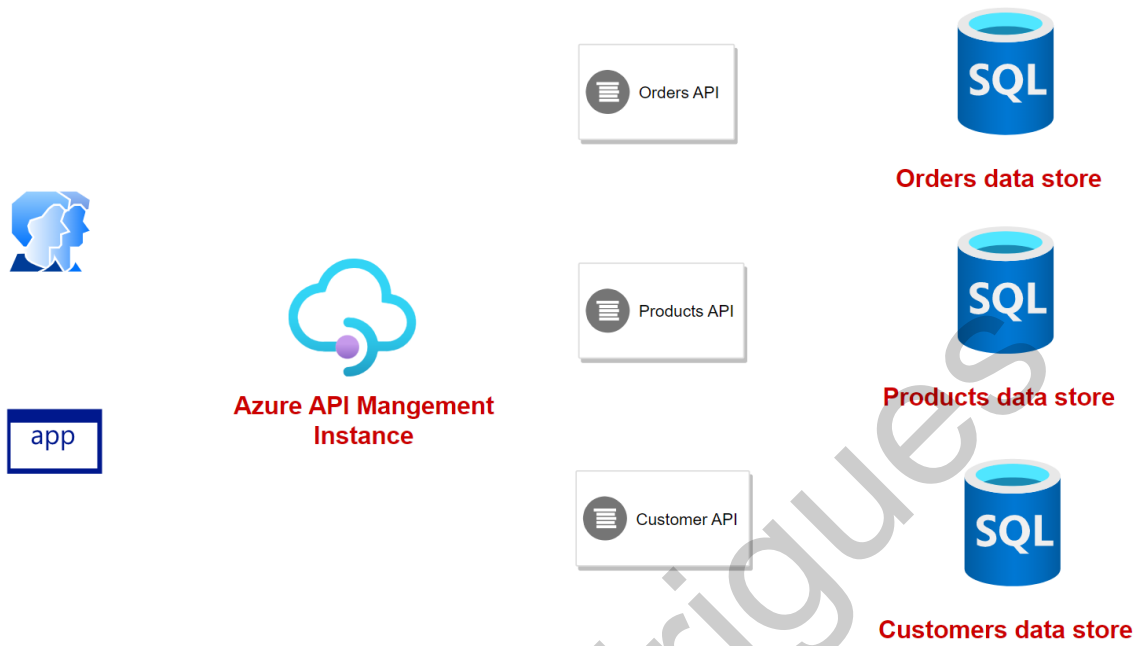


Lab- Azure Logic Apps



Azure API Management Service

Azure API Management

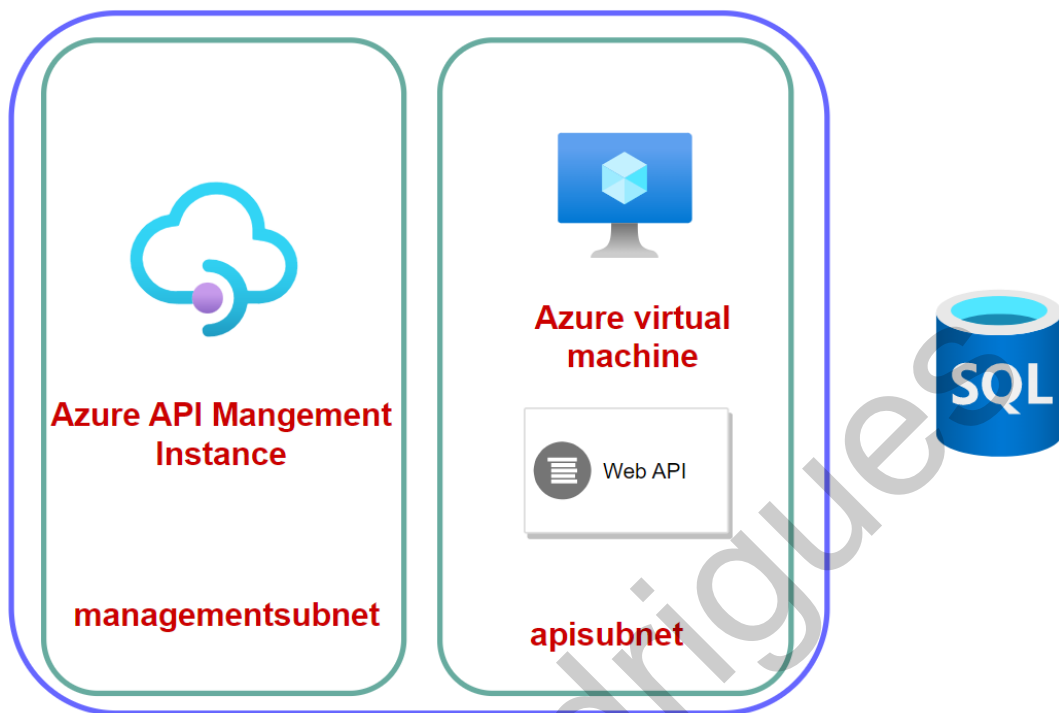


Allow customers to securely access backend API's

Abstract the implementation of the backend API's

Alan Rodrigues

Azure virtual network



Alan Rodrigues

Review of ARM templates



Azure virtual network



Azure virtual machine



Azure virtual machine



Azure Availability set



Azure SQL database



Azure storage account

You define your infrastructure as code

Create an Azure Resource Manager template

This is a JavaScript Object Notation file that actually contains the definition of the infrastructure

You can store the ARM templates in your source code repository along with your application code

Migration Patterns

Migration Patterns



Cloud Adoption Framework

Microsoft Cloud Adoption Framework



Your business wants to move applications to the cloud

Define strategy

- 1. Why do you want to move to the cloud
- 2. Would you be able reduce costs by moving to the cloud
- 3. Want to use cloud-native technologies

Plan

- 1. What are we going to need to move to the cloud
- 2. Can we move the application to the cloud



TB's of data , may need to consider how to move the data

The application is calling an external service.

data

Maybe because of legal purposes , we may not be able to move the data to the cloud

Ready

- 1. Keep our cloud estate ready
- 2. Design - What cloud services are we going to use

Adopt

- 1. Migrate your application
- 2. Enhance by building new cloud-native components

Secure

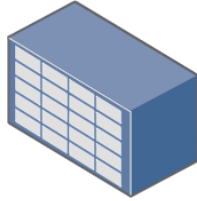
Protect your resources

Manage


Manage the operations

Data transfer options to Azure Storage

Data transfer options



Data center storage



Azure Storage accounts

How much data do you want to transfer

Do you have enough bandwidth to transfer files over the Internet

Is it a one time transfer of data?

One time large transfer of data

Import Export utility

Azure Data Box

Frequent transfer of delta data

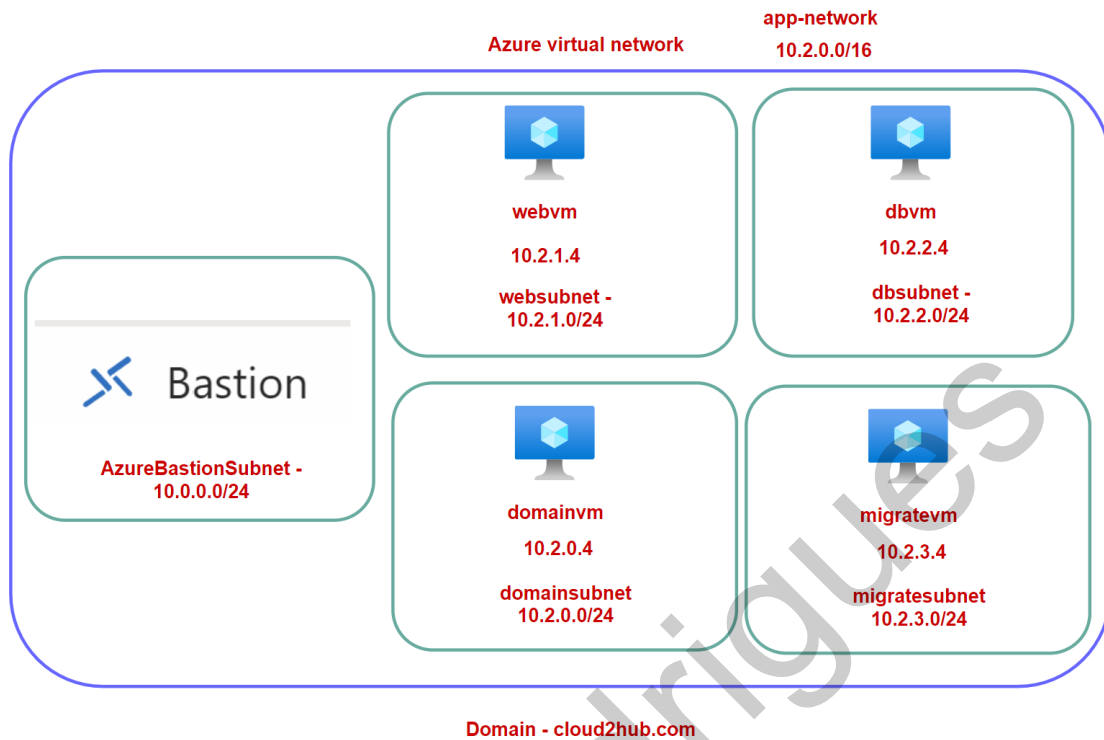
Azure storage explorer

AzCopy tool

Azure Data Factory

Exploring Azure Migrate- Infrastructure Readiness

Our Architecture



Alan Rodrigues