Unofficial Study Guide for AZ-305 Exam: “Designing Azure Infrastructure Solutions”

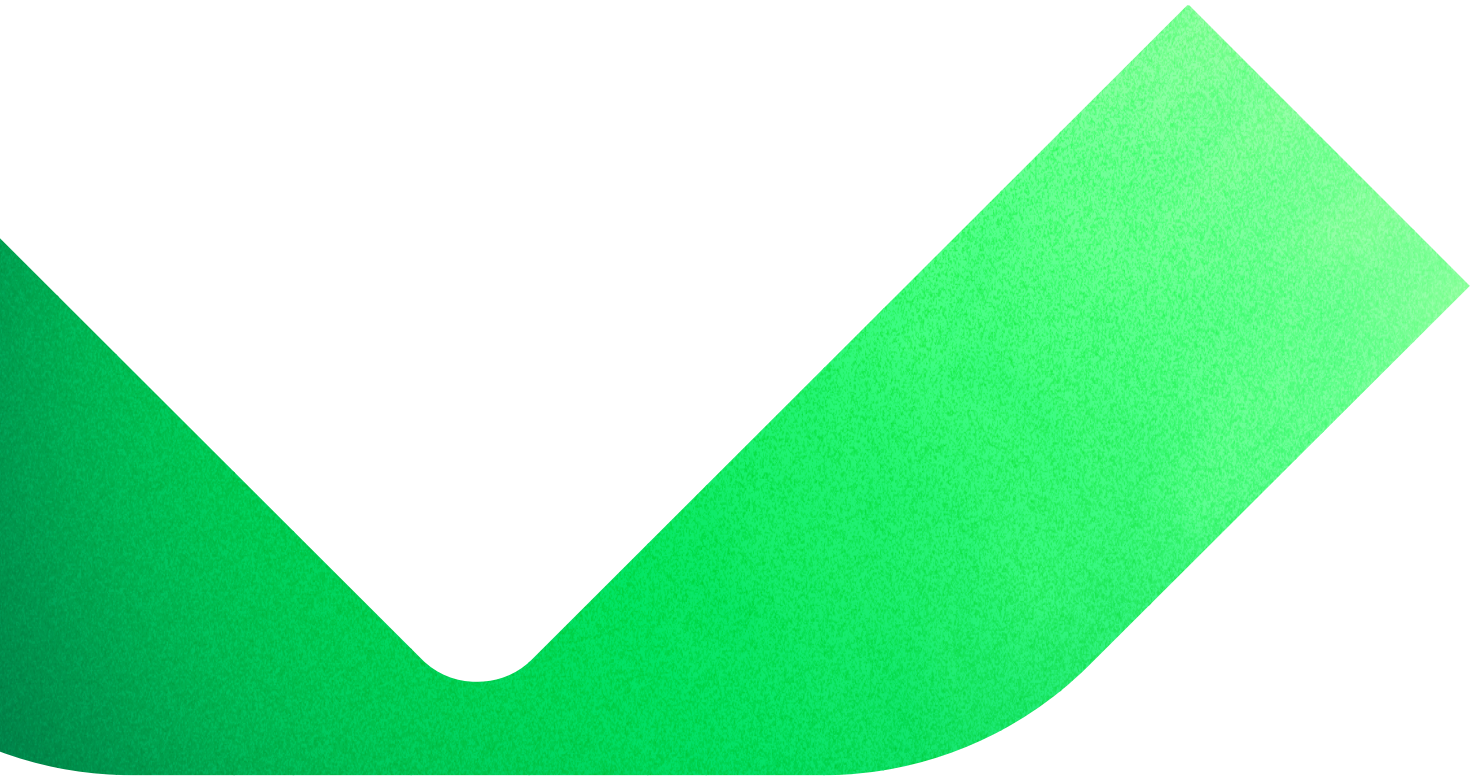


Table of Contents

Core concepts of Azure	4
Data structure and architecture	4
Governance	5
Compliance	5
Integration	5
Migration	6
Managing data loss and downtime	7
AZ-305: Designing Microsoft Azure Infrastructure Solutions exam: focus areas	9
Design identity, governance and monitoring solutions	12
Governance	12
Logging and monitoring	16
Authentication & authorization	17
Design data storage solutions	20
Data storage solutions for relational data	20
Azure databases	20
Data storage solutions for non-relational data	26
Data integration	28
Data path design strategy — hot, warm and cold	29
Design business continuity solutions	30
Backup and disaster recovery	30
Backup and recovery considerations	31
Azure Backup	32

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Recovery for blobs, files and Azure SQL	33
Azure Site Recovery	34
High availability	35
Design infrastructure solutions	38
Compute	38
Network solutions	47
Application architecture	51
Design solutions using the Azure Well-Architected Framework	54
Cost optimization	55
Operational excellence	56
Performance efficiency	57
Reliability	58
Security	58
Conclusion	61
About the Author	62

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

AZ-305, also known as the "Designing Azure Infrastructure Solutions" exam, is part of the Azure Solutions Architect Expert certification path. It evaluates your expertise in designing solutions that run in Azure as well as hybrid cloud environments.

In this white paper, we will first review some key aspects regarding Azure and background material for the exam pertaining to data loss and downtime. We will then delve into the different assessment areas of the exam and detail the topics you should familiarize yourself with to prepare for it.

This exam guide can be used by both new and experienced Azure professionals to study for the AZ-305 exam along with other official preparation materials. For cloud experts, this will serve as a refresher on core cloud concepts to springboard you into the data protection field. For infrastructure teams, this will relate core data management concepts you're already familiar with to specific Azure services.

Core concepts of Azure

Azure is one of the most preferred cloud platforms since it is backed by Microsoft's decades of expertise in managing an enterprise application landscape. Azure offers multiple services and tools — everything an organization requires to build, deploy and manage the end-to-end lifecycle of applications in the cloud. It offers enterprise-grade scalability, security, availability and data management capabilities.

Data resiliency in Azure is a hot topic, as organizations typically host mission-critical applications and data on the platform. As with any other cloud service provider, Azure also expects customers to implement certain guardrails to ensure data resiliency.

This is a journey for most organizations, especially those who are new to cloud since it involves building competencies in areas such as governance, compliance, data architectures, data structure integration and migration. All of these are contributing factors to data resiliency and protection, and we will briefly touch upon what you would be expected to do in each of these areas.

Data structure and architecture

Azure provides multiple data services that cater to different data storage and management needs. These include databases such as Azure SQL, Azure Cosmos DB, Azure MySQL, Azure Database for MariaDB and Azure Cache for Redis to name a few. Azure Storage is the object storage service available for unstructured data, while for data integration requirements, you can use Azure Data Factory.

Large-scale data storage and analytics mandate services such as Azure Data Lake. And then there are a plethora of additional relevant services just a click away in the Azure Marketplace.

Identifying the best-fit cloud services is the first step in building the right data architecture. Most will have resiliency built into them out of the box. For example, Azure Storage keeps three locally redundant copies of the data by default. However, if you need additional resiliency and protection, you will need to select advanced options such as globally redundant storage.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Note: There could be additional fees involved for all services once you incorporate additional data resiliency options into the architecture. For example, you can balance cost and resiliency by building out lower environments such as dev/test without duplicating resources and through measures such as backing up data.

Governance

There are multiple built-in capabilities that customers can use for data governance. Role-based access control (RBAC) helps guarantee that only specific users within the organization have access to Azure services. You can also use Azure policies to implement restrictions over the specific configurations of resources — for example, deployment geographies — allowed resource SKUs and naming conventions.

Azure Monitor provides complete monitoring, tracking and proactive management capabilities, enabling organizations to ensure proper governance of enterprise deployments. It can collect telemetry data from different sources such as Azure databases, storage, virtual machines and disks. The Logs Analytics service collects logs from various sources and analyzes them to detect access patterns, discover anomalies and implement the necessary controls.

Compliance

Implementing methods to adhere to legal obligations, security standards and corporate policies within the Azure environment is required to ensure compliance with Azure. Some standards such as ISO 22301 (Business continuity management systems) and System and Organization Controls (SOC 2) have specific controls organizations must implement for data resiliency in cloud environments.

The Azure platform's unique compliance requirements are handled by Microsoft, and audit reports are published in the Microsoft service trust portal.

Under the shared responsibility model of the cloud, users are responsible for managing the compliance of their hosted applications as well as data. For example, if data needs to be stored in two different regions for business continuity, it is up to the organization to configure the appropriate settings in Azure for multi-region availability. Similarly, for data resiliency requirements, you should use native or third-party backup solutions to meet any compliance controls specific to your environment.

Integration

Data migration and integration is the most crucial aspect of any cloud adoption journey. You should consider aspects such as the security of data in transit and at rest during a migration. Data integration is required when you need to combine and transform data from different systems to synchronize or consolidate it for further processing.

Some of the Azure services organizations can implement for integration are Azure Service Bus, Logic Apps, Event Hubs, Synapse Analytics and Data Factory. Each of these serves a unique function in the data architecture.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

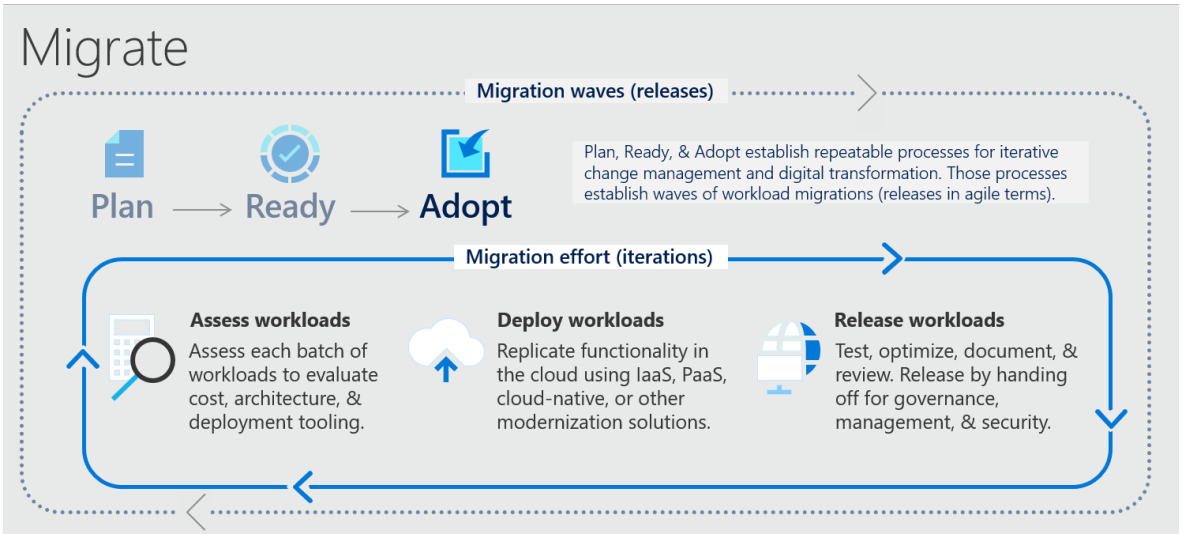
For example, Data Factory creates data pipelines that can extract, transform and load (ETL) data from multiple sources to target Azure services. While designing the architecture, you might have to consider regional redundancy to ensure data resiliency. In the same way, each integration service will have its own data resiliency considerations.

Migration

Data migration can be handled by services such as Azure Database Migration Service, Azure Data Migration Service and Azure File Sync. You should select the service that best suits the source data type and the Azure service to which you will be migrating data. You can also consider reputed third-party service providers for backing up and restoring data to the cloud if it suits your migration needs.

Cloud Adoption Framework

The Microsoft Cloud Adoption Framework for Azure provides prescriptive guidance on the tools, best practices and services that can be leveraged for a successful migration. This broadly includes activities in three phases: assess, deploy and release.



Microsoft Cloud Adoption Framework (Image courtesy: Microsoft)

The best practices guideline for users' reference includes an Azure migration guide overview, Azure migration best practices checklist, Cloud Adoption Framework migration model and a few sample scenarios based on defined methodologies. The prescriptive guidance provided by the framework can be leveraged by organizations to plan for different phases of their cloud adoption.

Migration options

While moving workloads to Azure, you can opt to go for lift & shift migration or refactor applications to use cloud native services.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Lift and shift

This type of application migration would require options similar to the source environment. This is often the short-term plan for organizations moving to cloud, especially if the priority is a data center exit. Organizations would not have the time and resources available to re-engineer their applications to use modern cloud services such as PaaS, FaaS or DBaaS.

Refactoring applications

Whenever this is a possibility, it is recommended that you go for it and leverage cloud-optimized services as much as possible for your architecture. The simple reason for this is the ease of use, the flexibility and the continuous improvements that can be expected for these services. Cloud-optimized services have baked-in features that help you make the most of all the features of cloud, such as scalability, high availability and security.

Migration tools

Azure Migrate assists with the different phases of cloud migration mentioned in the Cloud Adoption Framework. It is delivered as a unified service from the Azure Portal and has tools available for the discovery and assessment of workloads, migration, database migration, web app migration, etc.

It simplifies the migration of Windows and Linux workloads, SQL servers, Java apps, PHP apps, SAP HANA and specialized compute. If your application landscape consists primarily of these types of workloads, you should consider incorporating Azure Migrate into your cloud adoption plan.

Offline and online data migration

Azure Storage Migration Service can be used for the migration of unstructured data on-premises to the cloud. It can inventory your on-premises servers to get information on files and configurations. It is also suitable for transferring data from the identified source servers to the cloud; once the data is migrated, it can initiate a cut over to the target servers in the cloud.

Another useful service is **Azure File Sync**, which helps to centrally manage your organization's file shares using Azure files. The data in file shares are cached in servers closer to where the data is accessed, thereby improving the access speed and user experience. The data in on-premises file shares are kept in sync with Azure files and can help with lift-and-shift migrations of file share data to the cloud.

Managing data loss and downtime

While designing Azure environments, we should be aware of the fact that hybrid cloud deployments are the norm rather than the exception today. Hence, we need to create architectures that help manage data loss and minimize downtime.

Hybrid cloud data management has its own complexities associated with it, be it data integration, governance, security, performance management or resiliency. Out of all these, cloud resiliency is of paramount importance since downtime and data loss can lead to business loss and dissatisfied customers — something any organization seeks to avoid at any cost.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Although cloud offers several advanced data management and security capabilities, multiple factors can still lead to data loss and downtime.

Insider threats

Insider threats are caused by those who are authorized to access business-critical data and misuse their permissions with malicious intent. This can lead to data loss, data corruption, data mishandling or the injection of malicious code. Often, insider threats are not easily detected since they do not immediately flag abnormalities; this is primarily because a threat from inside an organization can access data using authorized credentials.

Note: These individuals do not always have malicious intent; their credentials can be compromised through phishing or social engineering attacks that lead to cyberattacks and data breaches. This could lead to cyber emergencies that many organizations are not prepared to handle.

Insider threats can also occur due to a lack of proper access controls for third-party vendors/contractors, resulting in the misuse or unauthorized access of data. Organizations frequently end up assigning more privileges than required to users, which can amplify the blast radius if their credentials are compromised. The principle of least privilege should be baked into every aspect of access control management to prevent this issue.

Accidental deletion of data

Human error or hardware/software failure are often behind data loss.

Users can delete data or corrupt it via incorrect commands or improper selection of files during a write operation. This frequently happens when users are in a hurry or lack familiarity with the system, especially if they are new. Hardware failures, such as disk corruption or machine crashes, can cause data corruption or make the data inaccessible. Similarly, software updates or application glitches can also lead to a data breach.

In some cases, the data can be recovered. In the worst-case scenario, there may be permanent data loss unless organizations regularly back up and store their data redundantly in multiple locations.

Ransomware and other cyberattacks

In the case of a ransomware attack, a hacker gains unauthorized access to data and encrypts it. Organizations then have to pay a hefty ransom to recover the data. Malware can also infiltrate your network through malicious websites, email attachments or compromised software, resulting in the corruption or deletion of files.

Another method of attack is a distributed denial-of-service (DDoS), where the system is overwhelmed with a flood of requests, thereby making it unresponsive to legitimate data access requests.

The ideal solution is to have protective measures in place to prevent the attack in the first place; however, proper data backup and recovery mechanisms can also help to recover the data from the most recent good copy.

The AZ-305 exam focuses on best practices that can be followed from the ground up to build Azure architectures that are resilient to such data loss risks.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

AZ-305: Designing Microsoft Azure Infrastructure Solutions exam: focus areas

The AZ-305 test is focused on evaluating the candidate's ability to develop various Azure solution components, such as infrastructure, security, data storage, networking and identity management. In the exam, subjects including Azure resource management, hybrid identity, virtual networking, data integration, security and compliance are covered. In addition to a solid understanding of Azure services, candidates are also expected to have a command of Azure best practices, architecture framework and patterns.

Listed below are the different areas covered by the exam and their weighting:

- Design identity, governance and monitoring solutions (25–30%).
- Design data storage solutions (25–30%).
- Design business continuity solutions (10–15%).
- Design infrastructure solutions (25–30%).

Your knowledge of these topics will be evaluated based on the modules presented for each below.

Design identity, governance and monitoring solutions (25–30%)

Design a solution for logging and monitoring

You should be able to differentiate between different types of log routing options available in Azure and recommend the best option. You should also have the appropriate skills to recommend monitoring tools and appropriate levels of logging.

Design authentication and authorization solutions

In this area, you will be evaluated on your knowledge of how to secure Azure resources using role-based access control (RBAC). You should also be able to recommend a solution for identity management and options for securing identities.

Design governance

Creating the right organization and hierarchical structure for Azure resources forms the basis of a secure landing zone, and you should be able to define this as per governance best practices. Your knowledge of Azure governance capabilities and your ability to recommend a solution for enforcing this will be evaluated in the exam.

Design identity and access for applications

You should be able to design solutions with appropriate access control for applications accessing other Azure services. This should factor in aspects such as user consent solutions, password and secrets management and integration with Azure Active Directory.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Design data storage solutions (25–30%)

Design a data storage solution for relational and non-relational data

For applications that use relational databases, you should be able to design the architecture using the right database type, size and tier. The scalability and resiliency of the service should be considered in the design. Data security settings such as encryption of data at rest, in transit and in use should be taken care of as well.

For non-relational data solutions, you will be evaluated on your ability to recommend services that balance features, cost and performance. Data protection, durability and access control mechanisms will need to be incorporated into the solution design.

Recommend a data storage solution

You will be evaluated on your skills in recommending the right-fit data solution for storing relational, non-relational and semi-structured data.

Design data integration

You should be able to design a solution for data integration and analysis that can fit specific use case requirements.

Design business continuity solutions (10–15%)

Design a solution for backup and disaster recovery

You will be evaluated on your ability to deliver a recovery solution for meeting the defined recovery time objective (RTO), recovery point objective (RPO) and recovery level objective (RLO) of workloads hosted in Azure, as well as in on-premises or hybrid environments.

You should also be able to provide the best-fit solution for backup and recovery services such as containers, compute services, databases and unstructured data.

Design for high availability

You should have the appropriate skills to evaluate the high availability requirement of applications and recommend solutions for compute services as well as relational/non-relational databases.

Design infrastructure solutions (25–30%)

Design a compute solution

You will be evaluated for your skills in recommending the right solution and architecture for compute requirements based on virtual machines, containers and serverless solutions. The compute solution will need to be right sized based on the workload requirements.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Design an application architecture

You should be able to recommend the right caching solutions, event-driven/messaging architecture, application management and automated deployment solution for your workloads. You should also consider appropriate API integration options for interacting with internal as well as third-party applications.

Design migrations

Cloud migrations should be designed in alignment with the Microsoft Cloud Adoption Framework for Azure. This involves an assessment of on-premises servers, data and applications that need to be migrated to the cloud. This should be followed by identifying the right solutions for the migration of applications, virtual machines, databases and unstructured data.

Design network solutions

You will be evaluated on your ability to design and implement a network architecture based on workload requirements. Applications should be optimized for network security and performance. You will need to factor in other considerations such as load balancing, routing and internet and hybrid connectivity based on design best practices.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Design identity, governance and monitoring solutions

To pass the AZ 305 exam, you should have a good understanding of the right governance frameworks, authentication and authorization solutions and the required monitoring controls for your Azure environments.

Governance

Data governance plays an important role in making sure that policies and restrictions are enforced per Azure best practices as well as specific organizational requirements. Azure offers several native services to implement this:

- Management groups.
- Subscriptions.
- Resource groups.
- Resource tags.
- Role-based access control.
- Azure policies.
- Azure landing zones.
- Azure Blueprints.

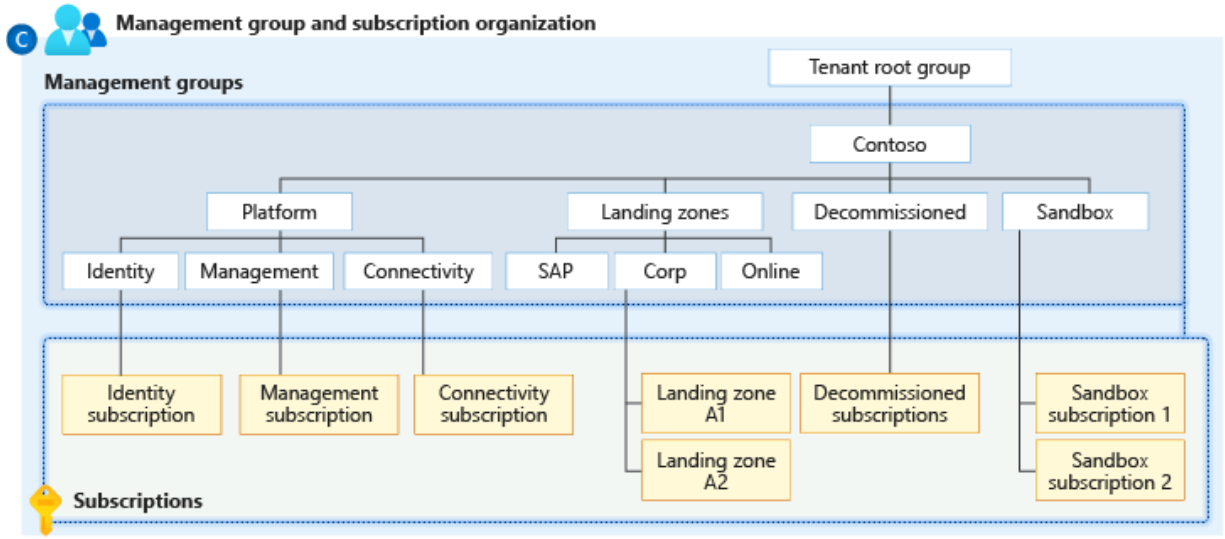
In the following sections, we will discuss each of these in turn.

Management groups

Azure recommends a hierarchical structure for your organization to apply proper governance policies at the right level. At the highest level of this hierarchical structure is a tenant root group that can have multiple management groups under it. These subsidiary elements group together multiple subscriptions and help administrators manage the policies, access and compliance for these subscriptions. Subscriptions act as logical containers and have their own billing for the individual Azure resources you create.

When you have similar policies and restrictions that need to be applied to multiple subscriptions, it makes sense to group them into a management group. There can be a maximum of six levels in a management group tree.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.



Management group hierarchy (Source: Microsoft)

Given the construct of management groups, the following are best practice considerations for logical organization:

- One management group can also have additional management groups as child nodes in the hierarchy. However, the best practice is to keep the group hierarchy reasonably flat to reduce management overhead.
- You can use a top-level management group followed by other management groups based on organizational/departmental structure, for example, finance, sales, IT etc.
- If your resources are spread across multiple geographies, you can consider a structure to implement relevant policies across different regions.
- If you have an application- or product-specific policy, access control restriction requirements or applications handling sensitive information, separate management groups can also be considered.
- Use a sandbox to experiment with new services. The policies and restrictions for sandbox environments do not need to be the same as those in other environments of the organization.

Subscriptions

Subscriptions are created in management groups in the resource hierarchy. You can group subscriptions that need to inherit the same policy and access control restrictions under one management group and apply these restrictions to the management group. Each subscription will have specific resource limits and quotas associated with it; hence you should consider the scaling requirements of your application and spread resources across multiple subscriptions if required.

Azure policies and Azure role-based access controls (RBAC) can also be applied at the subscription level. You can leverage this feature if you need to apply any additional policies to resources in a subscription for regulatory or compliance purposes. As subscriptions act as management boundaries, you can also take this into consideration when designing the hierarchy. For example, you can group resources that need to be managed by the same set of administrators into a single subscription.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Resource groups

Resource groups act as containers for deploying and managing Azure resources within a subscription. It is recommended to group resources that have the same lifecycle, as well as same policy and access control requirements, into a single resource group. For instance, the test, dev, QA and production resources of an application can be deployed in different resource groups.

All Azure resources should belong to a single resource group. The mapping here is 1:1, i.e., resources cannot be part of multiple resource groups, although a resource group can have multiple Azure resources. Note: A resource group will have its own region where the resource group metadata is stored; however, it can contain resources from multiple regions.

Resource tags

To categorize and organize Azure resources, resource groups and subscriptions based on specific attributes, you can assign to them key-value pairs known as Azure resource tags. This allows you to then filter and manage resources based on these tags.

For example, you can easily filter resources in a test environment or track costs incurred by different resources via tags in your Azure bill. You can add tags to map resources to different departments for this purpose.

Here are some of the best practices to consider for resource tags:

- When assigning tags for resources, consider your organization's naming convention. A tag can either be IT-aligned or business-aligned.
- Tags applied to a resource group are not automatically inherited by Azure resources in that group. Hence you can use Azure policies to assign one or more tags to resources when they are provisioned.
- You can create cost-related tags for resources to keep a track of budgets, configure alerts, track ROI, cost optimization, etc.
- Tags can be created based on mission criticality to streamline operations management and to keep track of SLAs.
- Tags can be created based on data classification so as to easily identify and determine security impact.
- You can create similar tags for resources associated with the same application. This helps with analysis of resources, identifying usage patterns and enable ongoing optimization.

Role-based access control

Fine-grained access control for Azure resources is achieved via Azure role-based access control (RBAC). Built-in roles, such as contributor, owner and reader, with specific permissions that can be assigned to resources. You can also create your own custom roles that have permissions aligned to your organization's requirements.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

RBAC roles can be assigned to users, groups or applications at a specific scope, i.e., either resources, resource group or subscription. These assignments are inherited by child resources depending on where they are applied. For instance, any resource that is part of a resource group will inherit the roles assigned to that group. This simplifies the access management process since you don't have to manage individual role assignments.

Make sure to follow the principle of least privilege when determining roles. Custom roles can be leveraged for more fine-grained access control in scenarios where pre-built roles might provide broader (riskier) permissions. Highly privileged roles, such as owner and contributor roles, should be assigned only to a limited number of users, and these should be regularly reviewed and audited. It is recommended that you assign roles to groups rather than users to help minimize the operational overhead of managing permissions, especially in large-scale deployments.

Azure policies

Azure policies help implement governance and compliance by enabling you to define and apply a set of rules and conditions across resources. You can apply either individual policies or a group of policies called initiatives. Azure policies are inherited by resources in the organization's resource hierarchy. They can be used to remediate non-compliant resources or prevent the provisioning of non-compliant resources.

Azure comes with its own pre-built policies, or you can create your own custom policies. Examples of restrictions that you can implement using Azure's own policies are:

- Restrict the type of VM SKUs that can be deployed.
- Audit VMs without disaster recovery enabled.
- Restrict the deployment of resources to specific Azure regions.
- Enforce tagging.

A full list of Azure's readily available policies can be found here.

The Azure compliance dashboard provides a bird's-eye view of an environment's compliance status against applied policies. You can initiate bulk or automatic remediation of resources from the dashboard.

It is important to understand the difference between Azure RBAC and Azure policies, as well as how both can be leveraged together to implement resource governance. **Azure policies ensure compliance of the resource state, while Azure RBAC implements controls over user actions on the resource.**

Azure Blueprints

Azure Blueprints bundles together Azure environment artifacts for implementing consistent governance and best practice patterns in a repeatable manner. This could include a set of resource groups, Azure policies, RBAC controls and Azure resource manager templates with security and compliance controls baked into them per your organization's requirements.

You can use Blueprints to create a resource group to deploy a set of defined resources; a blueprint acts as a scope for the policies and RBAC assignments. You can also include ARM templates in the blueprint to deploy resources even for complex environments. e.g., two-tier or three-tier applications. ARM templates are JSON files that define configuration of resources to be deployed in an environment. ARM templates help deploy resources as per defined organizational guidelines. The policies and RBAC definitions included in the blueprint will help implement the necessary access control and compliance guardrails.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

With Azure Blueprints, you can keep track of and audit the difference between what was actually deployed and what was intended to be deployed in terms of the relationship between the blueprint description and the assignment.

Azure landing zones

Azure landing zones define an environment that is built in accordance with best practices across the following design areas:

- Proper design, creation and deployment of Azure AD tenant and billing configuration.
- Implementation of security boundaries using identity and access management controls.
- Foundational network and connectivity design aligned with best practices.
- Resource organization through organization hierarchy design, i.e., management groups, subscriptions and resource groups.
- Identification of security requirements and the right security guardrails across workloads in the platform.
- Management baseline for ongoing operations that enable visibility, protection, recovery capabilities and operational compliance.
- Automation and DevOps tools for deploying landing zones and other resources.
- Regular auditing and enforcement of governance policies across environments.

Organizations should always deploy landing zones through code. You can use an iterative approach to make changes as the environment evolves and track the process.

Azure's landing zone accelerator implements a conceptual architecture with best practices pre-configured. This can be a good starting point for organizations, especially for greenfield deployments.

Logging and monitoring

Proactive monitoring is crucial to avoid downtimes, ensure systems availability and fine-tune performance. You can leverage solutions such as Azure Monitor, Log Analytics, Azure Workbooks and Insights to achieve this.

Azure Monitor

Azure Monitor is the built-in solution from Azure that collects and analyzes logs and metrics from different sources for monitoring purposes. It collects monitoring data from applications, operating systems, Azure resources, subscriptions, tenants and various custom sources.

Azure Monitor logs provide centralized storage for logs, as well as a built-in log analytics tool. You have the flexibility to define the parameters for collecting logs and how they are organized in the platform. Log Analytics is the primary tool for creating log queries in Azure Monitor. The outcome of the queries can be visualized in charts for identifying trends/patterns and configuring alerts.

Azure Monitor can collect metrics from Azure platform resources natively as well as from applications and virtual machines through additional agents. Collection takes place at defined intervals, with the data then stored in a time-series database. This helps provide valuable insights into system performance at different points in time.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

You can implement metrics for near real-time monitoring and alerting to flag critical events. The metrics explorer tool monitors and analyzes various resource metrics interactively to identify issues such as peak usage and resource constraints.

Azure Workbooks

Azure Workbooks is a feature of Azure Monitor that visualizes data collected from multiple sources. You can use this to drill down and analyze point-in-time information and create rich visual reports of the findings. Workbooks provide a holistic view of different aspects such as availability, usage, performance and health of various connected system components.

Azure insights

Azure Monitor provides visibility on the health of various resources in Azure such as containers, applications, virtual machines and Azure Cosmos DB. For example, Application Insights provides an in-depth view through the application management service (APM) that is part of Azure Monitor. Network Insights offers a topological view of network resources and their health as well.

You can use Azure insights and data analysis to create a customized monitoring experience for workloads in your Azure environments.

Authentication & authorization

Authentication and authorization mechanisms essentially hold the key to your kingdom, as unauthorized access to enterprise resources can compromise your entire organization. Azure offers comprehensive identity and access management solutions to handle this, including Azure Active Directory, Conditional Access and Identity Protection.

Identity and access management

In order to provide a seamless user experience, the ideal identity and access management (IAM) solution should have unified identity management, safe and adaptive access control and streamlined governance.

The identity service can either be native, such as Azure AD, or integrated with third-party identity providers through Security Assertion Markup Language (SAML) or OpenID Connect (OIDC) protocols. You can also leverage Azure AD business-to-business (B2B) or Azure AD business-to-customer (B2C) based on how your applications interact with third-party business applications and customers.

Azure Active Directory

Azure AD is the native IAM solution in Azure. You can implement this multi-tenant directory service either as a cloud-only identity management solution or extend it to your on-premises Active Directory for a hybrid solution. This is achieved through Azure AD Connect or Azure AD Connect cloud sync, which helps to integrate on-premises identities with Azure and provide a unified identity management solution. Accounts in Azure AD can additionally be used for assigning RBAC permissions to Azure resources.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Organizations should implement a unified IAM solution by combining their on-premises and cloud directories to limit operational overhead. The best option is to use a single Azure AD instance for your organization. While synchronizing on-premises accounts to Azure AD, it is recommended to exclude highly privileged accounts to mitigate the risk of attack vectors infiltrating the cloud from on-premises. Single sign-on (SSO) also lets users avoid having to remember and use multiple passwords.

Conditional Access

Azure AD Conditional Access helps in implementing granular access control based on the location of the user, the device from which the access is requested, trusted IP ranges and other factors. Conditional access helps allow, deny or enforce multi-factor authentication (MFA) based on the context of access; for example, customers who try to access applications using legacy protocols can be blocked. This allows organizations to grant access to applications only from managed devices. Access requests from untrusted devices or from unexpected locations that could indicate an infiltration will be denied.

There is additionally a report-only mode available in Conditional Access that helps you evaluate the impact of policies before implementing them. Similarly, MFA can be enabled for select users.

Identity Protection

The Identity Protection tool in Azure helps to detect and mitigate identity-based attacks. It calculates the real-time risk associated with a log-in attempt based on defined identity protection policies. If a risk is detected, the user will be required to use MFA, change their password or even get blocked.

Identity Protection helps prevent attacks that result due to compromised or leaked credentials.

Service principles

Service principals are implemented when a user or application requests access to an Azure resource that uses Azure AD for IAM. You can create and manage them using the Azure CLI, Azure PowerShell, Microsoft Graph and similar tools. You can then assign access policies and permissions to service principals as you would for a user account. Service principles can be used for application authentication and authorization in the same way as other user accounts.

Service principals are generally mapped to application objects, i.e., the definition of an application in Azure. An application can have multiple service principals mapped to it. You should always leverage service principals when authenticating external apps to Azure resources.

Azure Key Vault

Managing encryption keys, certificates and application secrets is always a challenge, especially in large-scale enterprise deployments. Azure Key Vault offers simplified management of your application keys, secrets and certificates. It helps to consolidate all app secrets in one place and ensure their encryption at rest and in transit.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Azure Key Vault has a standard tier that encrypts data using a software key and a premium tier that encrypts data using a hardware security module (HSM) protected key. In addition to enabling the centralized storage of secrets, Azure Key Vault also helps control access to them through access restrictions and policies. The key management process is isolated from the application, serving as a security boundary.

One of the recommended best practices while using Key Vault is to use separate key vaults for different applications. This helps reduce the impact in the event of a breach. You can fortify security using access policies, the principle of least privilege, virtual network service endpoints and firewalls.

To prevent accidental deletion and purging of keys, secrets and certificates in Key Vault, make sure to turn on the soft delete and purge protection service in Key Vault.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Design data storage solutions

How data is stored in your Azure environment depends on multiple factors such as data structure, data access patterns, integration requirements and security. Azure has multiple data solutions available for large-scale structured, semi-structured and unstructured data storage, and most organizations have a combination of these three data types:

- **Structured data** is stored in relational format with shared schema. The most commonly used form of structured data is in databases, where data is stored in tables that have columns, rows and keys.
- **Semi-structured data** doesn't fit into a table, row or column format, but it still has some structure; examples here include HTML, JSON and XML files.
- **Unstructured data** consists of non-relational data such as photos, audio/video files and documents that are not related but might be required to be stored together.

The AZ-305 exam will test your knowledge in identifying the right storage solutions for these data types. You will also have to know how to implement the design best practices for data storage in terms of access control, durability and data protection.

Data storage solutions for relational data

Relational databases are used in the backend of enterprise applications, where data is structured in tables with rows, columns and keys. While designing for relational databases, you should take into consideration the existing data landscape, its capacity and configurations if you are migrating from on-premises or from a different cloud. You should also take into account future growth requirements.

Azure databases

You have different options for databases, ranging from fully platform-managed to customer-managed solutions. The fully managed or PaaS (platform-as-a-service) options offer a hands-free approach to using databases once you select the right configuration for the database.

However, there are scenarios where you would need access to configurations that are not available on a PaaS platform. In that case, you should opt for an IaaS (infrastructure-as-a-service) approach, where the database can be deployed on an Azure virtual machine from the ground up.

Let's look at some of the different database solutions in Azure that fit into the above requirements, along with the features and caveats organizations need to consider.

Azure SQL Database

Azure SQL Database delivers the popular Microsoft SQL Server in Azure packaged as a PaaS service. A more refined category of the service is database as a service (DBaaS). The OS and SQL Server instances are abstracted so that customers can benefit from a fully managed service. The platform takes care of various operational and design aspects such as database tuning, high availability and backups.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

As it is a managed service, the platform automatically takes care of periodic updates to the latest version of the database. This significantly reduces the operational overhead associated with maintenance and upgrades, especially for organizations with large SQL data estates.

Azure SQL Server can either be deployed as a single database or as an elastic pool. Both are hosted on standard hardware platforms deployed, managed and maintained by Microsoft.

Single database

This type of deployment has a single database deployed on a logical SQL server with its own set of administrative controls. It will have a dedicated database engine and its own service tier, with a compute size that determines the resources available for the database.

In the Hyperscale service tier of the database, the compute and storage resources can scale independently as per application requirements. In the Serverless tier, the compute resources can be scaled and paused according to workload demand. The single database deployment is optimized for cloud-native application development use cases.

Elastic pool

The elastic pool deployment of workloads will consist of a group of databases that share the resources managed by a logical SQL server. This service is best suited for multi-tenant SaaS application patterns, where provisioning a single database per customer may not be a scalable option, especially if the performance requirements are not predictable.

Elastic pool provides a cost-effective solution, where you can purchase resources in a pool per a set budget, which can be shared by multiple databases. When one database instance needs additional resources based on its usage pattern, it can consume more resources from the pool based on defined parameters. Thus, databases with lower as well as high consumption can coexist in the pool and leverage the allocated resources.

Pricing considerations

Another important factor when selecting database resources is the pricing option. There are two available options: Database Transaction Unit (DTU) and Virtual Core (vCore).

- DTU is a purchase option that includes a pre-configured bundle of compute, storage and I/O resources.
- vCore gives customers greater flexibility in choosing the number of virtual cores required for the SQL database.

Azure SQL Database is designed to support large-scale deployments with up to 100 TB of data. Consider the following when designing SQL databases:

- If the resource requirements are simple and predictable, use the DTU pricing option with preconfigured resources.
- If you have existing investments in SQL licenses, the cloud cost can be reduced by leveraging the benefits of Azure SQL Server hybrid and vCore pricing.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

- If your resource usage is unpredictable, you can use elastic pools with vCore pricing and configure storage and compute resources for the pool. vCore pricing can also be used independently if you choose to go with multiple independent SQL databases.
- A serverless database will be a good choice for single SQL database configurations, as you will be billed only for the compute resources used.

Note: You should opt for Azure SQL Database is best suited when you need your applications to benefit from the features of the latest stable version of SQL Server.

Azure SQL Managed Instances

Azure SQL Managed Instances also offer SQL as a fully managed PaaS service. Unlike Azure SQL, this service offers additional control and configuration options to customers such as full virtual network integration, linked servers, Common Language Runtime (CLR), cross-database transactions, database mail and database virtualizations.

If your application architecture demands fine-tuning of instance-scoped features, then you can opt for Azure SQL managed services. It also offers full compatibility with on-premises SQL server deployments to facilitate migrations as well as hybrid deployments.

While designing Azure SQL Managed Instances, configure the maximum storage and CPU for the instances based on your application landscape requirements. Note that all the databases in the managed instance share the allocated resources.

Pricing considerations

Azure SQL Managed Instances use vCore-based pricing, where you pay a fixed price per hour for a combination of predefined vCore, memory and storage. You can also provision and use additional storage on a pay-as-you-go basis.

Note: Azure SQL Managed Instances are best suited for those customers with significant existing SQL investments and who want to migrate to Azure without having to redesign their application. They can still benefit from managed service benefits, including high availability, automated patching, automatic backups and minimal operational overhead.

SQL Server on Azure Virtual Machines

For full flexibility in an SQL deployment, be it a version of the server or its configurations (as is the case with on-premises), you can opt to deploy an SQL server in a virtual machine. These deployments follow the infrastructure-as-a-service (IaaS) approach. You can deploy SQL servers on Windows as well as Linux virtual machines.

With an SQL server deployment on a virtual machine, customers have full administrative control over how they deploy, configure and manage the SQL database. However, this also adds additional responsibilities such as patch and OS updates on the administrators. You should also consider factors such as high availability, scalability, backup frequency, DR configurations, etc. Hence, it is best to choose this option only if you have manpower available for managing the SQL servers.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Pricing considerations

Customers opting to deploy an SQL server on a virtual machine can leverage Azure Hybrid Benefit and reuse their existing Windows and SQL server licenses to save on costs. You can either use the available SQL server images in Azure, which includes the licensing cost, or go for a bring-your-own-license (BYOL) approach.

If using SQL server images provided by the platform, you will be charged on a per-minute basis for the license, along with the compute charges. In a BYOL model, only the compute charges will apply.

Note: SQL deployments on VMs are best suited if you are migrating applications that would need OS-level access and fine-tuning of SQL configurations. It is recommended for lift-and-shift migration scenarios, where the priority is to move to the cloud without delay and with minimal or no changes.

Azure Cosmos DB

Azure Cosmos DB is a fully managed, NoSQL relational database service that supports multiple database APIs. This includes leading open-source databases like MongoDB, Apache Cassandra, NoSQL and PostgreSQL as well as Azure Table. The choice of API to be configured depends on the application architecture requirements. You can use Azure Cosmos DB with greenfield deployments or for migrating data from any of the supported databases.

Azure Cosmos DB is a multi-region distributed data service with assured SLAs and single-digit millisecond response times. It is recommended for modern app development, as it eliminates the operational overhead while addressing key design requirements such as automated scaling, patch updates and high availability.

Customers can easily migrate applications written for Azure Table storage to Azure Cosmos DB with minimal changes since they use the same data model. When compared to Azure Table, Azure Cosmos DB has much less latency, typically less than 10 ms for reads from over 30+ regions. It also provides automated indexing, which provides faster query response times. You can choose from different consistency levels depending on how much your application can compromise on availability, throughput, consistency and latency.

Pricing considerations

Unlike Table storage, which provides only consumption-based pricing, Azure Cosmos DB provides both provisioned capacity as well as consumption-based pricing.

Database scalability considerations

Azure SQL Database offers dynamic scalability, where the compute, memory, I/O and storage can be scaled to fit application usage patterns, without compromising on application performance or availability.

With a single database deployment, you can choose the DTU or vCore pricing model and define the maximum available resources for each database. With elastic pool, you can configure the resource limits for all databases created in the pool.

Database scaling can either be done vertically or horizontally.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

With vertical scaling, i.e., scaling up or down, you increase the compute resources available for the database using elastic pools. Note that a pool needs to be configured with resources required to cater to peak usage, even if the average usage remains low.

Horizontal scaling entails adding or removing databases to cater to performance or capacity requirements. Also known as scaling out, horizontal scaling can be implemented using database sharding or read scale-out provisioning. With sharding, larger databases are split into smaller components, including a primary read-write replica and multiple read-only replicas for improved performance and simplified management. With read-heavy applications, it is helpful to offload reads to a read-only replica, which is known as a read scale-out configuration.

Database availability considerations

Azure SQL availability is dependent on the service tier, which in turn decides the managed SQL deployment architecture. vCore service tiers include General Purpose, Business Critical or Hyperscale. DTU service tiers can be Standard/Basic or Premium.

vCore General Purpose and DTU Standard/Basic tiers are best suited for deployments that need a balance of storage and compute options. Resiliency is managed via Azure Service Fabric, which manages the failover of instances to nodes with storage capacity, similar to how SQL Server's failover cluster instances (FCI) work.

vCore Business Critical and DTU Premium tiers should be considered for applications that need high resiliency and low latency, including mission-critical business applications. The architecture used in this tier is similar to SQL Server's Always On availability groups in the backend. It offers three secondary replicas, of which one can be used as a read-only replica free of charge.

The vCore Hyperscale tier should be used when you need highly scalable storage, especially for read operations. It can support large databases, up to 100 TB, and uses a tiered architecture with caches and page servers that enable faster read operations at scale. You can also take snapshots of databases for instantaneous backups and restores.

Data security at rest, in motion and in process

Protecting data stored in database services should ensure the security of data at rest, in motion and in process. Data-at-rest protection means that data being persisted to the storage layer should be secure. When data is in motion, i.e., during data transfers between public or private endpoints, it should also be secured so that it is not intercepted or manipulated. The same goes for data in process, i.e., while being used by various applications.

Transparent data encryption (TDE) for data at rest

Azure SQL uses TDE to encrypt and decrypt databases, their backups and log files in real time. As the encryption is transparent, it does not require any change at the application layer. An entire database is encrypted using a database encryption key (DEK), which can either be service-managed or customer-managed.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

SSL/TLS encryption for data in motion

Azure SQL Database encrypts all data communications between the client and server through SSL/TLS encryption for all connections. This is applicable to Azure SQL as well as Azure SQL managed databases. You can also use a VPN to encrypt the traffic in hybrid deployments. Azure ExpressRoute is another option that can be considered in hybrid architectures, where the traffic flows through a dedicated WAN link.

Dynamic data masking for data in use

For protecting data being processed by applications, you can use dynamic data masking. In simple terms, it masks sensitive data in database query results without altering the data in the database. This allows you to prevent unauthorized access to sensitive business data, in accordance with the principle of least privilege. Data masking policies can be configured through different options, such as Azure Portal, PowerShell and Azure Rest API to name a few.

Database considerations summary

To summarize, consider the following aspects while selecting databases for your Azure Architecture:

Scalability

- Use vertical scaling as per usage to increase the compute resources available for the database using elastic pools.
- Implement horizontal scaling using database sharding or read scale-out provisioning.
- Split larger databases into smaller components through sharding.
- Offload reads to a read-only replica, which is known as a read scale-out configuration.

Availability

- Choose Azure SQL service tiers based on your availability requirements.
- vCore General Purpose and DTU Standard/Basic tiers are best suited for deployments that need a balance of storage and compute options.
- The vCore Hyperscale tier should be used when you need highly scalable storage, especially for read operations.

Security

- Azure SQL uses TDE to encrypt and decrypt databases, their backups and log files in real time.
- Azure SQL Database & Azure SQL managed database encrypts all data communications between the client and server through SSL/TLS encryption for all connections.
- Use a VPN to encrypt the traffic in hybrid deployments.
- For protecting data being processed by applications, you can use dynamic data masking.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Data storage solutions for non-relational data

Now that we have explored the services for relational data, let's look into services that can be used for non-relational data: Azure Storage, Azure Files and Azure managed disks.

Azure Storage

Azure Storage caters to unstructured data storage requirements and combines different storage services like blobs, files, tables and queues. The features and capabilities of Azure Storage depend on the type of storage account, which is determined by its location, cost, replication capabilities, etc.

The different types of Azure Storage accounts are as follows:

- **Standard general purpose V2:** used for blobs, queues, tables, files and disks stored as page blogs; the most common type of general-purpose storage, as it suits a wide variety of use cases.
- **Premium block blobs:** support blob storage only (block blobs or append blobs); can also be used for data lake storage; recommended for applications that need scalability and high transaction rates.
- **Premium file shares:** used to host managed file shares.
- **Premium page blobs:** support page blobs, i.e., index-based sparse data structures used in virtual machine disks.

While designing for storage, make sure you choose the right location, especially if your organization has specific data residency requirements. Storage cost is another matter of consideration, and it is directly dependent on the data size as well as redundancy configuration.

You can either opt for locally redundant, zone-redundant or geographically redundant storage accounts. Redundancy configurations ensure that data is protected against local, zonal and regional outages:

- **Locally redundant storage (LRS):** three copies of data are stored locally in the same data center. This avoids a single point of failure but is not sufficient to protect data if the entire data center goes down.
- **Zone-redundant storage (ZRS):** data is replicated to data centers in three different Azure availability zones in the same region. This configuration protects data from zonal failures, as the data would still be accessible from the available zones in the region.
- **Geo-redundant (GRS) and geo-zone-redundant storage (GZRS):** there are three additional copies of data in a paired secondary region, replicated asynchronously, to protect your data from regional failures. GZRS provides additional resiliency against availability zone failures as well as in the primary region.

Azure files

This file share service is fully managed and hosted in Azure. Cloud files can be accessed through standard file share access protocols such as SMB and NFS as well as via the Azure Files REST API.

You can consider Azure Files in your architecture for replacing on-premises file share services like network-attached storage (NAS). It can also be used for lift-and-shift migrations of legacy applications that have a dependency on file shares.

The services can be configured and accessed in a few clicks rather than setting up an entire file server architecture from scratch using virtual machines.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

You can configure the file shares to be locally redundant, zone-redundant or globally redundant based on your application resiliency requirements. You can choose from standard or premium file share accounts depending on your performance needs. Standard file shares offer low-cost file shares backed by HDD, which can be used for standard workloads with low performance requirements. For improved latency, premium storage with high IOPS and bandwidth backed by SSDs can be used.

Azure managed disks

Virtual machines use Azure managed disks for static data storage. Disks can be provisioned by specifying the size and type of the disk, while the process is managed by the platform in the backend.

There are four types of managed disks:

- **Standard HDD:** uses HDD in the backend; recommended for non-critical data such as tests or backups.
- **Standard SSD:** uses SSD; recommended for enterprise applications that are not heavily used or for test and development purposes.
- **Premium SSD:** uses SSD and offers higher throughput than standard SSD; recommended for business-critical production applications and well as applications with high performance requirements.
- **Ultra-disk:** uses SSD and offers the maximum IOPS of all the disk types; recommended for transaction-heavy and latency-sensitive applications such as databases or SAP HANA

You will need to perform a cost performance analysis to identify the disk type best suited for your workloads. For increased performance, you can also consider disk caching options to optimize read/write operations.

Storage security

Azure offers several controls for ensuring the security of storage accounts. However, aligning with the cloud shared responsibility model, it is up to the customer to identify and configure the storage security controls per their organization's requirements.

- In general, you can consider the following best practices for Azure storage security:
- Adhere to Azure's set of guidelines called "Azure security baseline for Storage." These recommendations can be monitored using Azure Defender in order to take corrective action where required.
- Ensure the principle of least privilege for accessing storage accounts through the use of RBAC and shared access signatures.
- Configure secure transfer for your storage accounts so that only secure connections are accepted and all other requests are rejected.
- Encrypt data at rest in Azure Storage as well as Azure managed disks either using Microsoft-managed (default option) or customer-managed keys stored in Azure Key Vault.
- Implement network layer security through features such as firewall policies and rules, virtual network service endpoints and private endpoints.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Data storage solutions for non-relational data

During cloud migration, it is inevitable that you will need to plan for large-scale data ingestion, processing and analysis. With ever-expanding data estates, this data could be in scales of terabytes or petabytes that would be too complex to be handled by traditional databases. Hence, you would need data-integration solutions that can bring together data from multiple sources, analyze it and deliver critical business insights.

AZ-305 will test your understanding of data integration solutions that can help with this process, such as Azure Data Factory, Azure Databricks, Azure Synapse Analytics and Azure Stream Analytics.

Azure Data Factory

Azure Data Factory creates workflows that can help with large-scale data movement and transformation, integration of data from multiple sources and extract, transform and load (ETL). Ingested data can be stored in Azure Blob, and Azure Synapse Analytics can be used for further processing.

Azure Data Factory helps with low-code/no-code data integration through an authoring and monitoring tool available in Azure Portal. It is offered in a fully managed serverless architecture, which eliminates the overhead of infrastructure management.

With support for 90+ connectors, you can use Azure Data Factory for data integration from multiple sources, either in the cloud or on-premises.

Integration and analytics with Azure Databricks

Azure Databricks is a managed big data and machine learning platform based on Apache Spark. Data engineering and machine learning teams with expertise in Apache Spark can easily start using this service. It also has built-in APIs to support common languages such as Python, Scala, R and Java.

You can use Azure Databricks to support multiple data integration and machine learning use cases. For example, the integrated machine learning environment can be used for training new models, developing new features and tracking experiments. You can also use specific data clusters for targeted jobs, which data scientists can then further analyze.

Teams can collaborate using Databricks' built-in shared workspaces to manage the data science lifecycle, while machine learning experts can build recommendation engines. It is also suited for other use cases like intrusion detection.

Azure Synapse Analytics

Azure Synapse Analytics delivers a packaged solution for data integration, storage and analytics. You can use the service to query data at scale, even in a serverless architecture, and to process both relational and non-relational data. It also assists with descriptive analytics, diagnostics and predictive as well as prescriptive data analysis.

Azure Synapse Analytics helps with in-memory data analysis, which offers better performance for certain data sets. It can be used for both batch and stream data types and enables real-time analytics of streaming data, which is very helpful for applications that need access to immediate insights.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Azure Synapse Analytics fits well into hybrid architectures since it supports integration with on-premises data warehouses.

Azure Stream Analytics

Azure Stream Analytics is a fully managed low-latency, highly available streaming data processing service capable of processing millions of complex events per second. It also supports the analysis of both structured and unstructured data.

The input data streams can be generated from sources such as IoT devices, sensors or social media feeds. The data from these streams is analyzed using the event processing capability of the service to generate actionable insights, identify trends, etc. The outcomes of the analysis can be sent to Power BI for visualization.

You can also use Azure Stream Analytics to send data to downstream services such as Azure functions or Azure Queue Storage. For low-latency requirements, you can run the service on Azure Stack Edge or IOT Edge.

Data path design strategy — hot, warm and cold

Cloud handles data differently compared to how data is stored and managed on-premises. Because of this, you need to consider the strategies for using the right services for hot, warm and cold data paths.

Warm data path

With a warm data path approach, data analysis happens in near-real time as the data flows through the cloud services. Azure Stream Analytics would be an apt solution for analyzing warm data. For storage of warm data, consider database services such as Azure Cosmos DB or Azure SQL.

Cold data path

In a cold data path, analysis and pattern identification occur over a period of time. You might also need to use historical data or integrate with a warm data path to deliver a holistic view. The cold data path requires long-term data storage, which can be accomplished using services like Azure Blob Storage, Files, Data Lake Storage, etc.

Hot data path

This refers to real-time data analysis and processing. The data is often latency-sensitive and also needs quick analysis for faster outcomes. Low-latency and high-throughput storage and analytics services should be considered here.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Design business continuity solutions

Downtime and data loss are a costly affair, as they can lead to loss of customer trust and financial penalties. Backup and recovery solutions are a critical part of any Azure architecture design to ensure business continuity. There are multiple services available in Azure for implementing an effective backup and disaster recovery (BDR) strategy.

Backup and disaster recovery

When organizations move or deploy mission-critical workloads to the cloud, one of the major considerations is making sure applications remain available and are resilient to downtime. As with any other hosting environment, cloud may also experience downtimes, so it is important to integrate services into the architecture that can perform granular point-in-time recovery of data. In rare cases where the entire cloud region/zone is affected, you also need to think about possibly bringing the application back online in a different region or zone. Let us take a look at the core concepts of backup and DR, before exploring the solutions offered by Azure in this space.

RPO and RTO

RPO defines the maximum data loss that can be tolerated by an organization in the event of a disaster, while RTO defines the maximum tolerable downtime.

For example, an RPO of one hour means the organization can tolerate one hour of data loss without significant impact. An RTO of one hour means the system should be up and running within one hour of an outage. The choice of backup solution will vary depending on how high or low the RPO and RTO are. A low RPO and high RTO would require a backup solution, but a high RPO and low RTO would need a recovery solution based on continuous replication.

High availability vs. disaster recovery

High availability is about designing architecture to minimize downtime and make the components resilient to failure. This can be achieved through different options such as using redundant component, configuring automated failover or through load balancing.

Disaster recovery on the other hand helps recover your system from a major disaster, failure or cyberattack. For DR purpose you should have a copy of your data and application in a different isolated location so that you can bring back your applications online if the primary location becomes unavailable.

Backup and recovery considerations

While planning and designing a backup and recovery solution, you need to factor in different aspects of the architecture.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

System components

Most applications will have different components or tiers as part of their design. So, you need to ensure that backup and recovery is factored in for each of these. For example, a two-tier application with a VM for a web frontend and databases as a backend should have a disaster recovery (DR) plan for both tiers if it needs to be redeployed from the ground up in the event of a system failure.

Usage patterns

The usage patterns of a workload play a key role in deciding the recovery process. For example, you could be lenient with a development and test environment in terms of backup frequency, but that cannot be the case for a production environment with mission-critical data.

MTTR and MTBF

In system architecture, you must also factor in the acceptable mean time to recovery (MTTR) and mean time between failures (MTBF) to meet your organization's service level agreement (SLA) and redundancy requirements.

The value of MTTR, i.e., how long before a system recovers, is critical in defining the type of DR solution you implement: either active or passive.

MTBF, on the other hand, is the average time between consecutive failures of system components; this could be a determining factor for your backup frequency.

Snapshots

Snapshots help create point-in-time copies of data volumes. They help restore critical application data in the event of data corruption due to software/hardware failures, malware attacks, human error, etc.

All modern backup solutions can take incremental snapshots, where only data that has been modified since the last snapshot is captured. This helps reduce the time taken for the backup and also the storage space required for the backup copies.

The recovery of data from snapshots is faster and can help meet stringent RPO and RTO requirements.

The overall SLA of your application will depend on the SLAs of its components and how each of them is designed for resilience and DR. Azure offers various native services that architects can leverage to achieve even the most stringent SLA.

Azure Backup

Azure Backup is a cloud-based backup and recovery solution for end-to-end protection of your workloads on-premises as well as in the cloud. Using Azure Backup, you can ensure the security of your Azure VMs, files, blobs, SQL servers (in IaaS deployment) and SAP HANA databases in VMs.

For on-premises workloads, backups are performed via the Microsoft Azure Recovery Services (MARS) agent, Microsoft Azure Backup Server (MABS) agent or Data Protection Manager (DPM). The MARS agent helps with file, folder and system state backups, whereas the MABS can take backups of Hyper-V/VMware virtual machines as well as applications such as SQL, Exchange and SharePoint.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Some of the advantages of using the Azure Backup service are:

- Cloud-scale backup and recovery without having the overhead of building and maintaining a backup and recovery infrastructure.
- Free data transfer for the backup and recovery operations, without limiting inbound and outbound data.
- Centralized management of the backup solution with integrated monitoring and alerting capabilities.
- Ability to store data either on-premises for hybrid deployments or in the cloud for short-term/long-term retention.
- Various storage options such as LRS, ZRS and GRS for redundancy of backup copies.
- Application-consistent backups to ensure faster recovery of applications to a working state.
- Out-of-the-box integration with Azure services and configuration through a few simple steps.

Azure stores backup data in vaults. A vault is a storage entity that comes in two types: Backup Vault and Recovery Services Vault. Backup Vaults help with backups only and are used to store the backup copies of Azure disks, blobs and PostgreSQL servers. Recovery Services Vault can be used by both Azure Backup and Azure Site Recovery services for disaster recovery purposes. These support data sources, such as: Azure virtual machines, SQL servers, SAP HANA and file shares.

Since vaults store backup data that is crucial for business continuity, it is important to make sure that you have considered various critical aspects:

- Define backup policies that meet the RPO and RTO requirements of your system and enforce them across vaults in your subscription via Azure Policy.
- Use RBAC to ensure that only authorized administrators can access your backup vaults; follow the principle of least privilege so that backup data does not fall into the wrong hands.
- The right level of redundancy should be factored in either through LRS, ZRS or GRS; use LRS if local data protection will suffice, opt for ZRS to protect from zonal failures and go with GRS to protect from regional failures.
- Depending on your data segregation requirements, choose different vaults for Azure Backup and Azure Site Recovery; you can also have multiple vaults created if workloads are distributed across different subscriptions.

Recovery for blobs, files and Azure SQL

Blobs, files and Azure SQL are commonly used services in Azure architectures, so making sure the data stored in these services is protected against corruption or data loss is a must. Azure Backup offers out-of-the-box integration with Azure Blobs, files and Azure SQL.

Azure Blobs

Backups of Azure Blobs are stored in Azure storage accounts instead of Backup Vaults. This entails a continuous backup process, which can be used for point-in-time recovery of data. This is called an operational backup and is enabled at the storage level rather than centrally via the Azure backup. It uses capabilities such as soft-delete and blob versioning to protect data in the storage accounts.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

When configured, operational backups enable delete lock on your storage accounts to protect them from accidental deletion. This will also prevent unauthorized users with malicious intent from deleting a storage account.

You can configure backup policies for blobs with a maximum retention period of 365 days. The blob versioning capability helps you restore data to a working version, should there be an incident of data corruption or accidental deletion. With the soft delete feature, any deleted data is retained for a defined period of time; this is available for blobs as well as containers.

Azure files

Azure files can either be protected using share snapshots or Azure Backup for files. Share snapshots are supported only for SMB files and take point-in-time snapshots of Azure files. Snapshots are taken at the root level i.e., it will include all files and folders in that file share. During recovery, you can opt for granular individual file restore.

Snapshots are immutable, i.e., once created, they cannot be changed. You can, however, read, copy or delete them.

The recommended approach for protecting Azure Files is via Azure Backup integration. You can automate the snapshot backup process using Azure Backup and configure it for different retention periods i.e., daily, weekly, monthly, yearly, etc. You can also take on-demand backups when required, for example, before making an application code change. Azure Backup supports instant restore of files as well from the snapshots.

Azure SQL

Azure SQL and Azure SQL Managed Instance have automated backup enabled, which helps protect critical business data stored in these services:

- Full backups are taken every week. These include the database and transaction logs.
- Differential backups are triggered every 12-24 hours. With these, all changed data since the last full backup is captured.
- Transaction log backups are typically taken every five to10 minutes. This helps fail back to a backup copy of transaction logs in the event of failure or corruption.

Automated backups help with point-in-time recovery of databases to a known good state within the retention period. The restore is done in the same SQL server, and once the data is verified, you can delete the original database. You can also restore the database to a different geographical region. This can help you recover your databases, should the primary Azure region become unavailable.

Automatic backups have a default retention period of 35 days. However, if you need longer retention of backups for regulatory or compliance purposes, Azure SQL offers long-term retention (LTR), which allows you to keep backup copies for up to 10 years.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Azure Site Recovery

Azure Site Recovery is a managed cloud-based DR solution that can be integrated with hybrid and multi-cloud architectures. It delivers end-to-end orchestrated disaster recovery as a service (DRaaS), which helps meet your business continuity requirements.

The service replicates data between primary and secondary locations, where the primary location can either be on-premises, in Azure or in workloads hosted on-premises on physical servers or virtualized platforms such as VMware or Hyper-V. It also supports the replication of virtual machines hosted on AWS.

Azure Site Recovery (ASR) provides a centralized solution for DR configurations as well as for initiating the failover and failback of workloads. You can enable disaster recovery for your Azure VMs directly from Azure Portal through a few simple steps. For the protection of VMware machines, you need to deploy an additional replication appliance on-premises; this can be deployed as a VMware VM using an open virtualization format (OVF) template. Hyper-V machines can be replicated to Azure by installing the Hyper-V Site Recovery Provider and Recovery Services Agent on Hyper-V hosts.

If you have a Hyper-V virtual machine manager (VMM) for centralized management of Hyper-V hosts, you can install the site recovery provider in the VMM for DR orchestration. For legacy on-premises deployments with physical servers, you need to deploy an on-premises Azure Site Recovery replication appliance that can manage the disaster recovery configuration, failover and failback.

Application-aware recovery

The business continuity and disaster recovery (BCDR) service offered by Azure Site Recovery is application aware. It offers out-of-the-box integration with enterprise applications such as SQL Server, Active Directory, Oracle, SAP, SharePoint, Exchange, etc. These applications can be deployed either on physical or virtual machines (VMware and Hyper-V) on Windows/Linux operating systems. The service uses app-consistent snapshots in the backend to deliver a near-synchronous application-aware replication.

Recovery plans

You can also create recovery plans for your multi-tier application and integrate scripts for pre- and post-failover steps, which can bring down the overall RTO. Azure offers an automation library with scripts for leading enterprise applications that can be leveraged here, or you can use your own scripts tailored to your organizational requirements.

RPO & RTO

You can configure ASR and fine-tune your recovery plans to ensure that you meet your RPO and RTO requirements.

The VM recovery RTO SLA offered by Azure is two hours. In practical scenarios, the failover typically takes only a few minutes, although the RTO would still be dependent on other factors, including pre- or post-failover steps.

The RPO is dependent on the replication frequency, which is continuous for Azure/VMware VMs and 30 seconds for Hyper-V.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Failover and failback

You can perform planned or unplanned failovers using ASR. Planned failovers can be triggered for scheduled maintenance events, whereas unplanned failovers are performed during unexpected downtimes. Once the primary site is up and running, you can also trigger a failback for all environments, except for physical servers. For physical servers, the only option is to fail back to a VMware virtual machine.

Test failovers

It is recommended to test your BCDR plans using the ASR test failover feature. It helps to simulate failovers to an environment isolated from production without any disruptions so that data replication is not interrupted during the test failover process. You can periodically test your failover plans to make sure you can meet the defined RPO and RTO goals and that there are no surprises during an actual failover.

High availability

The high availability (HA) design considerations should align with business continuity and disaster recovery requirements. For incidents that affect the entire landscape of your enterprise, such as regional failures and data center failures, you need a broader DR strategy. HA helps recover your workloads from localized incidents including server failures, component failures, etc. In terms of RPO and RTO, this would most probably be minutes instead of hours.

High-availability architecture options will be different based on the type of service you are using and the delivery model, i.e., IaaS or PaaS. Most PaaS services, for instance, will have high availability features built into them. IaaS deployments using virtual machines, on the other hand, would need more involvement and planning from the administrators to ensure you have all the right guardrails in place to ensure application resiliency.

High availability for virtual machines

You can ensure high availability of your virtual machines in Azure by deploying them either in availability sets or availability zones.

Availability sets:

Availability sets eliminate single points of failure in your architecture within a data center. When you allocate machines to availability sets, they get deployed across different fault domains and update domains.

Machines in a fault domain will be connected to the same network and power source in a data center. Since machines are distributed across different fault domains, you can eliminate single points of physical infrastructure failure.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Update domains, on the other hand, indicate physical infrastructure components that could be rebooted at the same time, for example, during an update or maintenance activity. Distributing virtual machines across update domains ensures that your application is not impacted by such downtimes.

For multi-tier applications, it is recommended to deploy the application components in separate availability sets for additional resiliency.

Availability zones

While availability sets can protect against failures within a data center, you also need protection from events where the entire data center could be impacted by an outage. Availability zones can guarantee this protection.

Each supported region will have multiple availability zones, which will be a unique physical location within that region. There will be a low-latency network connection between the zones in a region. However, note that a single zone may not be restricted to a single data center.

Although deployments in availability zones protect from zonal failures or events where entire data centers are impacted, it introduces a minor latency of about one ms. It could, however, still support high-availability architectures like Always On availability groups used by SQL.

If you are opting for an SQL deployment on VMs, you could decide to have additional high-availability architectures, i.e., an Always On failover cluster instance (FCI), Always On availability group (AG) and log shipping. While using FCI, you should also use an internal load balancer so that users can connect to its unique name. With FCI, when one node of the SQL cluster fails, the instance is restarted on another node.

High availability for PaaS solutions

As mentioned earlier, PaaS services will have default high availability built into it. This takes away the administrative overhead of setting up additional servers and configuring them to achieve high availability. Hence, it is the recommended option for cloud applications, except for circumstances where you are, for example, migrating legacy solutions or deploying services that are not yet available as PaaS.

There are high-availability options for managed databases. For example, Azure SQL and Azure SQL DB managed instances have built-in auto-failover groups and active geo-replication out of the box. For Azure Database for MySQL, any hardware-level issues impacting the service are handled by the platform, where the built-in failover mechanism will protect you from downtime and also ensure data consistency. For node-level disruptions, the node is recreated to bring the service back online transparently.

At the application level, it is recommended to implement retry logic so that the user experience is not impacted during the database recovery. For databases, another recommended practice is to use a read replica offered by the PaaS service. This assists in offloading the read and reporting operations. When the read replica is deployed in a different region, it also helps with high availability.

Hybrid solutions

Hybrid cloud deployments are commonplace today, as most organizations are on a digital transformation journey to the cloud and want some part of their infrastructure to be deployed on-premises for reasons such as data residency, compliance, etc.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Some cloud-native services have high availability features that can be extended to on-premises for hybrid deployments. You can extend SQL Managed Instance to on-premises using the SQL Server transactional replication feature. However, you would most likely see IaaS deployments in a highly available hybrid deployment, meaning you would need to extend the Azure network to cloud for such implementations.

While there are multiple hybrid connectivity options available, it is recommended to use Azure ExpressRoute, as it offers dedicated bandwidth and connectivity to your on-premises environment. This will help meet the stringent RPO and RTO requirements for a highly available workload deployment.

One of the most common hybrid architectures is to use Azure as a DR site for workloads on-premises. This can be achieved through out-of-the-box solutions like Azure Site Recovery or native application capability. For example, you can deploy an SQL availability group configuration with a secondary replica in the cloud. You can also store backups to cloud in hybrid deployments for active-passive high availability; this allows you to recover your environments to the cloud using the backup data in the case of availability issues for your on-premises deployment.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Design infrastructure solutions

Infrastructure solutions form the backbone of any Azure architecture. They provide the necessary compute, network and application services, as well as define the framework for migrating workloads to the cloud. These components ensure the deployment, scalability and performance of applications hosted in the cloud.

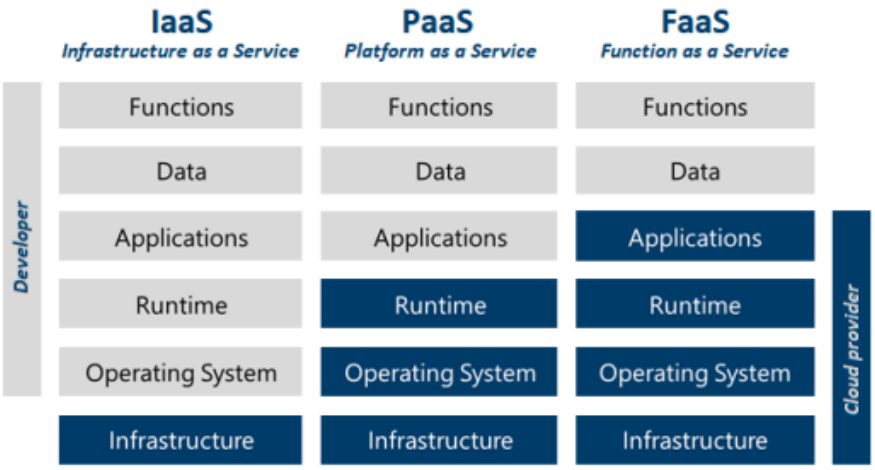
There are multiple services that can provide compute capacity for your workloads: Azure VMs, containers, App Service, Azure Kubernetes Service (AKS), Azure Functions, Logic Apps to name a few. Azure offers a comprehensive software-defined networking layer to ensure the connectivity of your application components within the cloud as well as to other hybrid environments. This includes virtual networks, gateways, firewalls and load balancers, everything required to enable this connectivity.

Following application architecture best practices is important, for example, selecting the right design patterns for messaging, caching and API integration while designing your Azure architecture. The tools and techniques you use for migrating your applications to the cloud will also help you avoid downtime and disruption. Let's take a look at the aspects that go into designing an infrastructure in detail.

Compute

The selection of compute resources is dependent on your deployment model, i.e., whether you are going to use IaaS, PaaS or a serverless/function-as-a-service (FaaS) approach. This will in turn depend on the type of application, its mode of operation and the level of configuration flexibility you need.

IaaS,PaaS and FaaS hosting options at high level can be represented as shown below:



Compute hosting options(image courtesy: Microsoft)

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Infrastructure as a Service (IaaS)

In this deployment model, you deploy applications in virtual machines provisioned and managed by you. You will need to handle aspects such as the associated storage, network connectivity, scalability, high availability, security, etc. This approach is required if you plan to run highly customized workloads.

Another scenario is if you have specific requirements like specialized hardware or GPUs where you might have to opt for IaaS. This deployment gives the same flexibility as managing your on-premises data center, except that the underlying hardware is owned and operated by Microsoft.

IaaS is often the recommended approach if you are looking for a quick migration to cloud from on-premises without refactoring your applications. Also, if you go for the IaaS model in Azure, you can deploy your workloads in virtual machines.

Platform as a Service (PaaS)

With PaaS, you get a pre-built hosting environment with all the required binaries and dependencies installed to run your application. It helps you focus on your application rather than worrying about managing the underlying infrastructure or keeping it up to date. Cloud-native applications are generally considered best suited for PaaS services. You can also use PaaS to migrate applications from on-premises if there is room for minor refactoring to suit the PaaS environment.

PaaS is much easier to administer and manage compared to IaaS. However, you will not be able to do OS-level configurations if you opt for PaaS services.

The commonly used Azure compute services for PaaS are Azure Batch, App Service, AKS and Container Instances.

Function as a Service (FaaS)

With FaaS, you can deploy your code (also known as a function) on the platform without having to set up the hosting environment for it. The code will be executed dynamically in an environment where the compute resources are completely abstracted from the user. The resources are also allocated as required by the code, and you will be charged only for that amount.

Functions are usually executed based on a trigger, for example, in response to a request in a queue or an HTTP request. This is well suited for event-driven application architectures. The services that offer FaaS in Azure are Azure Logic Apps and Azure Functions.

Microservices architecture

With more organizations opting for container-based microservices architecture, you need to select the Azure services best suited to host them. These decisions will depend on factors such as the source container orchestration platform, scalability and ingress/egress options. You might also want to consider containers for migrating commercial off-the-shelf (COTS) applications if they are available as pre-configured container images.

You can refer to the following decision tree from Microsoft to plan the most suitable compute service for your application. Note that the process will be either complex or easy depending on your specific application requirements:

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

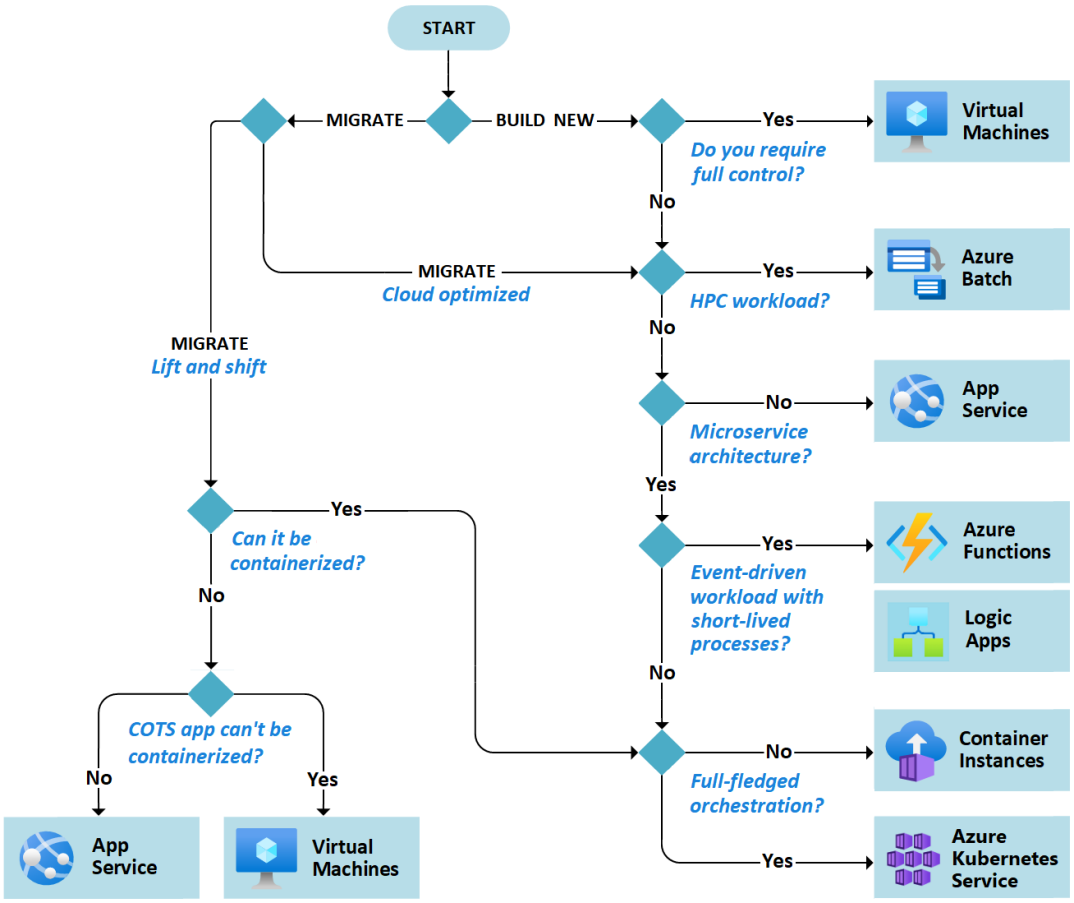


Figure 1: Decision tree for Azure compute service (Source: Microsoft)

Azure VMs

The fundamental compute service in Azure to deliver IaaS is Azure VMs. Virtual machines are the fastest and simplest way of migrating or deploying your workloads to the cloud. They offer on-demand, flexible and scalable compute options with end-to-end control of your application starting from the OS. The ownership of maintaining the physical platform for running the VMs is with Microsoft; however, everything starting from the VM size, location and network to security and settings is up to the user to decide and manage.

Virtual machines can be used for hosting production as well as test/development environments. If the application demands additional resources, you can enable horizontal scaling by adding more machines. By extending the on-premises network to Azure, you can also extend your data center to the cloud. VMs can be used in lift-and-shift cloud migration scenarios as well as greenfield application deployments.

Let's take a look into some of the design considerations when choosing virtual machines for your Azure architecture.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Location

You want to deploy your virtual machines in an Azure region that is close to the user. You may also want to select a specific region based on data residency or compliance requirements. In a multi-region architecture, you should deploy virtual machines in multiple regions, with traffic being distributed across them using services like Traffic Manager or Azure Front Door. The disks associated with the machine can also be stored in a region that you select.

While deciding on the location, you should also verify that the machine SKU you want to use is available in that region, as not all machine SKUs are available in all regions. The cost of machines could also be different depending on the region.

Virtual machine type

Azure offers different virtual machine SKUs, categorized as general purpose, compute optimized, memory optimized, storage optimized, GPU-based and high-performance compute. Depending on the nature of your workloads, you will need to select the appropriate machine SKU and size. The SKU that you select will decide the level of compute capacity, number of disks that can be attached and network bandwidth available for the hosted workloads; hence, this is one of the crucial design decisions of your Azure architecture.

General purpose SKUs offer a better balance between cost and performance; these can be used for lower-tier test/dev environments or for websites with minimal traffic. With a higher CPU-to-memory ratio, these machine types are best suited for low-to-medium traffic applications, web servers, etc.

Memory-optimized virtual machines, on the other hand, are configured with a higher memory-to-CPU ratio, which is required when it comes to, for example, cache-heavy applications and in-memory analytics. Storage-optimized virtual machines have higher storage throughput and I/O, i.e., database servers. There are other special purpose SKUs for specific workloads that need a GPU or high-performance computing.

Cost

Azure bills hourly for virtual machines. If a VM is not switched on for the entire hour, you will be charged for the minutes the VM was in use. The cost components will include the compute charges, the cost of the disks and the operating system charges if you are using an Azure-managed OS image.

You can configure an auto shutdown of the test/development machine during a specified time to save on compute charges; however, be aware that you would still have to pay for the managed disks associated with the VMs.

Disk performance

You need to select the disk size and type i.e., Standard/Premium/Ultra Disk, based on your application performance requirements. You can refer to the "Design data storage solutions" section in this eBook to better understand the different managed disk options available. The number of disks that can be added to a VM will depend on the VM SKU, so you need to plan according to the workload's capacity requirements.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

The managed disks are created and managed by the platform, meaning you do not need to worry about the backend configuration required for this.

Availability

You can deploy virtual machines in availability sets or availability zones for highly available deployments. With availability zones, Azure guarantees 99.99% availability with two or more instances deployed across availability zones. You can refer to the "Design business continuity solutions" section of this eBook for a deep dive on high availability configurations using availability sets and availability zones.

Scalability

For horizontal scalability, you can deploy virtual machines in Azure Virtual Machine Scale Sets and place them behind a load balancer. Virtual Machine Scale Sets help with centralized management of the scaling configuration based on various performance-related and other metrics. This helps ensure horizontal scalability for your workloads according to your application usage patterns.

Virtual machines in these scale sets can be deployed across multiple availability zones so as to ensure high availability in addition to scalability.

Operating system

Azure offers most of the leading Windows and Linux operating system images in a pay-as-you-go model that can be used to deploy virtual machines. There are third-party images available from independent software vendors (ISVs) in the Azure Marketplace with pre-installed software for popular application stacks, such as PHP, WordPress and LAMP.

Using these images helps accelerate the deployment timeline since the components come installed and configured. You can also bring your own 64-bit operating systems to the cloud by uploading the images to Azure Storage.

Automation

There are many infrastructure-as-code (IAC) tools available for automating the deployment and management of virtual machines. You can use Azure Resource Manager (ARM) templates, Bicep, Terraform or Ansible to achieve this.

You can also leverage custom script extensions to configure virtual machines after deployment. Azure supports cloud-init as well to customize Linux virtual machines post-deployment for tasks such as installing software packages and copying files.

Azure Container Instances

Containers help package your application and its dependencies in a lightweight image that can be deployed on supported containerization platforms. It uses the concept of OS virtualization, where you deploy multiple containers in the same operating system, and it makes them much faster, more portable and more efficient when compared to virtual machines.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

To deploy standalone containers in Azure, you can use Azure Container Instances. This is recommended for smaller applications, tasks or jobs that do not require a full-fledged container orchestration platform like Kubernetes. It is commonly used in Dev/test scenarios as well. Some of the benefits of using Container Instances are as follows:

- Given their lightweight nature, you can launch them in seconds rather than minutes, giving you faster access to applications.
- You can specify the resources (CPU and memory) required for your containers and ensure that you pay only for the compute resources that you're using.
- Azure Container Instances can be used to deploy both Linux and Windows containers.
- You can use Azure Files as persistent storage for your Azure Container Instances.
- Containers that share the same life cycle can be deployed in a container group, and they will share the same computer resources, network, storage, etc.

Deployment considerations for Azure Container Instances

You should use either Azure Container Registry or any other private registry such as Docker Trusted Registry to store your container images. It is not recommended to use public registries for your business application containers.

You can use built-in tools such as Azure Defender for container registries to ensure that your container images are tamperproof and protected from malicious activities.

Azure Policies will ensure the compliance of your container images and remediate any non-compliant images or even prevent the deployment of such images.

Lastly, make sure to monitor the resource usage pattern of your containers to keep an eye on possible performance bottlenecks. You should also monitor elements such as logs and file access patterns for security purposes. Built-in container monitoring capabilities of services such as Azure Monitor and Azure Container Insights can be leveraged for this purpose.

Azure App Service

Azure App Service is a platform-as-a-service option from Azure for deploying your web applications, batch jobs, mobile backends, etc. It supports multiple popular programming languages and frameworks, such as .NET, NET Core, Java, Node.js, PHP and Python, and you can deploy your applications in both Linux- and Windows-based environments. It has built-in features for traffic management, security and scalability that simplify your application deployment and management process.

Let's look at some of the features and considerations when using Azure App Service.

Deployment and management

Azure App Service environments are managed by the platform; this means that all patches related to the operating system and the frameworks are deployed and managed by the system so that developers can focus on the application. You can also deploy containers in App Service, where both Windows and Linux containers are supported.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

We can connect your Azure App Service to your enterprise applications such as Salesforce and SAP through out-of-the-box connectors and also with social media platforms. It also offers multiple options for authenticating users, for example, Azure Active Directory, Microsoft accounts, Google accounts or even social media accounts.

DevOps integration

Azure App Service also integrates with several DevOps services such as Azure DevOps, Bitbucket or GitHub, enabling the continuous integration and deployment of your applications. Docker Hub or Azure Container Registry will let you deploy containers to Azure App Service, after which you can integrate the service with your Azure DevOps deployments.

Depending on the App Service plan tier, there will be different deployment slots. For example, you can use a staging slot to deploy and test the application and then swap it to a production slot once it is ready to go live. You can also use the autoswap feature to automatically move the deployment to production if such a validation is not required.

Type of deployment

Using Azure App Services, you can deploy either web apps, API apps or WebJobs depending on your use case. For API apps, you can use the built-in Swagger support to publish your APIs in the Azure Marketplace, which can then be consumed by any HTTP or HTTPS client. You can also use Azure App Service to build your mobile apps, as it supports native iOS, Xamarin, Android, etc.

Pricing

Azure App Service offers multiple pricing tiers for you to choose from. You can either start with a free tier or a shared tier. With the free tier, you pay only for the compute resources your application consumes. In the shared tier, there is a quota of CPU miners that can be used by each app in the tier.

For higher environments, you can go for dedicated plans, such as Basic, Standard, Premium, Premium v2, Premium v3, etc., where the charge is incurred based on the number of VMs used for the app deployment and scaling.

Azure Kubernetes Service (AKS)

For standalone containers, you can use Azure Container Instances, but if you need extensive management and orchestration of containers, then the recommended solution is Azure Kubernetes Service (AKS). It offers the popular container orchestration service Kubernetes as a packaged managed solution that helps centrally manage and scale applications deployed using a microservices architecture.

The control plane of AKS is free, i.e., you only pay for the agent nodes in which your containers are deployed. AKS can be used for deploying new microservices-based applications or lifting and shifting existing containerized applications from on-premises. For example, if you have existing containerized applications deployed in an on-premises Kubernetes cluster, you can simply lift and shift the application to AKS.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

AKS provides customers with the flexibility of open-source Kubernetes while also allowing them the benefits of a managed service, such as autoscaling, automated updates and self-healing.

Deployment considerations for AKS

Depending on the scale requirements of your application, you can opt for two types of autoscaling in Azure Kubernetes Service: Horizontal Pod Autoscaler and Cluster Autoscaler. In Horizontal Pod Autoscaler, the number of pods is increased when the application usage increases, leading to higher resource usage. With Cluster Autoscaler, the number of nodes in the cluster is increased when there is not enough capacity in the existing nodes to deploy the scheduled pods.

Integrating Azure Kubernetes Service with Azure AD lets you leverage the existing Azure AD identities to authenticate to the Kubernetes cluster. This provides an integrated sign-in experience for your Azure users to applications hosted in AKS.

The Container Insights feature of Azure Monitor allows you to monitor the performance of AKS clusters as well as the hosted containers. Azure Kubernetes Service also offers integration with leading monitoring and visualization services such as Grafana and Prometheus.

AKS additionally provides cordoning and draining processes for upgrading your cluster nodes so that there is minimum disruption to the hosted applications.

If you're planning to host compute- or graphics-intensive workloads in Azure Kubernetes Service, you can leverage GPU-enabled node pools, which can be configured to use virtual machines with single or multiple GPUs.

For persistent volumes in AKS, you can use Azure disks or Azure Files. Azure disks are recommended for volumes that will be accessed by a single pod, while Azure files are suitable for shared/concurrent pod access.

You can integrate AKS with an existing virtual network in Azure so that the deployed applications can communicate with other resources connected to the network.

AKS supports Docker images and has out-of-box integration with Azure Container Registry, a secure managed private registry service from Azure.

Azure Functions

Azure Functions offers a serverless option for running your code in the cloud without managing the infrastructure. It supports code created in multiple languages such as C#, Java, JavaScript, Python, Go, etc. Azure Functions integrates with popular development tools such as Visual Studio, Maven and Visual Studio Code for direct cloud deployments and debugging.

You can also instrument the code to use Azure Application Insights for telemetry and analysis. The code can be configured to be executed in response to specific events. Azure Functions can also scale on demand based on usage.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Design considerations for Azure Functions

Azure Functions has a default timeout of 30 minutes for all plans other than consumption plans. For consumption plans, the timeout is 300 seconds. Hence, it is recommended not to use Azure Functions for long-running code. You can, however, possibly split up such long code into smaller functions and integrate them to work together.

A Durable Function, an extension of Azure Functions, can be used for the stateful execution of code. This is good for scenarios like function chaining when one function's output becomes the input for a subsequent function.

It is recommended to create function applications depending on resource requirements to meet performance and scaling targets. Also, functions should be written in such a way that they can manage exceptions and continue to operate even if there are failures with any associated services or resource limitations.

By default, Azure Functions resources are regional services. However, you can also deploy them in a multi-regional active/active pattern using Azure Front Door to distribute the requests between functions in different regions.

Lastly, it is recommended to allocate different storage accounts for functions, especially if they require higher IOPS.

Azure Logic Apps

Azure Logic Apps is part of Azure's integration services and helps connect services and applications in hybrid deployments. With Azure Logic Apps, you can create and automate workflows that can integrate different components of your application such as services, system, data, etc. The workflow signature logic apps can be created using either a graphical user interface or JSON files and then triggered based on events.

You can use pre-built connectors for Azure Logic Apps to access enterprise services and applications such as Salesforce, Twitter and Office 365. For example, you can create a logic app to respond to a specific event and send an email notification using Office 365.

Some things to consider when using Azure Logic Apps are:

- Include Azure Logic Apps in the design when you have multiple applications and systems that need to interface with one another.
- Use the Logic Apps execution engine to scale the application based on usage.
- Use conditional expressions such as switch statements, loops and Boolean expressions in your code to make data decisions and implement highly complex conditional logic.
- In addition to built-in characters, you can also create custom connectors for logic apps if you have an existing SOAP API or REST API.

It is not recommended to use Azure Logic Apps for real-time data processing or complex business rules-based requirements.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Azure Batch solutions

High-performance computing applications with large-scale parallel computing requirements can be effectively run in Azure using Azure Batch. Azure Batch lets you build and manage pools of compute nodes, plan jobs to run on the nodes and auto scale resources up or down as required.

Azure Batch also helps you optimize the speed and cost of your workloads, manage task dependencies and implement job restrictions. It can be used in multiple scenarios such as machine learning, rendering, simulation and testing.

Best practices for using Azure Batch

Pools consist of compute resources that are required to execute an Azure Batch job. If you have a small number of tasks in a job and the expected runtime is not long, it is not recommended to allocate a new pool for every new job, as this could lead to reduced runtime due to the additional time required for node allocation.

Since the pool is a security boundary in Azure Batch, you should create multiple pools if there is a security requirement to isolate jobs from each other.

You can group together multiple tasks into a single job and give it a unique name, which will help monitor the logging of Azure Batch activity. It is simpler to group tasks in this manner than create a sprawl of Batch jobs, which becomes very difficult to track.

The workflows you create should be resilient against failures, as Batch pools could experience downtime due to underlying hardware errors. In case of a node failure, the task will be automatically re-executed on a different compute node, so the Batch job should be designed to handle such retries.

You can deploy batches in different pools across different Azure regions so that the jobs are resilient against regional failures.

Network solutions

Designing a proper network solution is a crucial step in Azure architecture design. A network solution defines how the different components of an Azure architecture communicate with each other and with external resources. The network solution can also impact the performance, security, scalability and reliability of a workload deployed in Azure.

Azure networks should be designed in such a manner that they provide reliable and robust experiences for users accessing workloads from different locations.

Network architecture considerations

While designing your Azure architecture, you should keep in mind the functional and non-functional requirements of your workloads. You also need to consider the following aspects when designing a network solution:

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

- The connectivity options and network topology for Azure resources, including virtual networks, subnets, virtual network peering, VPN gateways, ExpressRoute circuits, etc.
- The network IP address spaces required and their further segmentation based on the connectivity and isolation needs of workloads.
- The network security and access control policies for Azure resources, including firewalls, service endpoints, application security groups, network security groups and private endpoints.
- Comprehensive monitoring and logging of the network using services such as Network Watcher, Network Performance Monitor and network security group flow logs.
- Network performance, load balancing and an optimization approach using a load balancer, traffic manager, application gateways, CDN, etc.

Best practices

Some of the best practices to follow while designing your virtual network and subnet design are as follows:

- Use different subnets for the different tiers of your application, each with a unique IP address range.
- Ensure there is no overlap between the IP address ranges of each of the subnets. This is especially important when you are connecting different networks through options such as virtual network peering or hybrid connectivity.
- Each Azure subscription has quotas and limits on the number of interfaces and IP addresses you can have in a virtual network; make sure you take this into consideration when designing your network topology.
- If you need to control the network traffic coming in and out of a virtual network, make sure that you are implementing network security groups that implement these restrictions.
- By default, Azure allows all resources within a virtual network to talk to each other, irrespective of whether they are deployed in different subnets. To change this, you need to override the default route with custom ones per your traffic routing requirements.

Common design patterns

Single virtual network

In this pattern, all the components of your applications are deployed to a single virtual network in the same region in Azure. Note that virtual networks are regional resources in Azure, i.e., they cannot span across regions; however, you can create multiple subnets inside a virtual network for different tiers.

To implement network traffic segmentation and ingress/egress traffic restrictions, you can use network security groups or application security groups. For example, Network Security Groups ensure that traffic between web- and application-tier subnets is allowed only through specific ports. Additional services such as Azure Firewall or network virtual appliances enable you to implement additional network traffic restrictions.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Multiple networks with peering

In this pattern, you can use multiple virtual networks in different Azure regions and connect them through the virtual network peering service. This is helpful in multi-region deployment patterns where your application is being accessed from multiple geographical regions.

Even though there can be a one:many relationship between virtual networks in this pattern, note that network peering is not transitive. For example, if network A is peered with Network B, and Network B is peered with Network C, resources in Network A and C cannot communicate with each other without a direct peering relationship.

Segmentation using network security groups, application security groups, network virtual appliances, Azure Firewall, etc. can be implemented in this pattern as well.

Hub-and-spoke topology

This is a more advanced network pattern where one virtual network acts as a connectivity hub with multiple spoke virtual networks connected to it through virtual network peering. You can implement centralized network segmentation and traffic control at the hub, which will usually have some kind of a network virtual appliance to do this, governing traffic flow through the topology.

This is also an easily expandable pattern when your network topology grows, as you can centrally manage the security posture at the hub.

Hybrid network connectivity

Organizations can link their on-premises infrastructure to the Azure cloud using one of the several hybrid connectivity options provided by Azure. These allow for seamless data flow and integration between resources hosted on-premises and resources hosted in Azure. Below, we present the primary hybrid connectivity options in Azure.

Azure VPN

You can establish site-to-site VPN connections using Azure VPN Gateway between an Azure virtual network and your on-premises network. Communication is secured using IPsec/IKE protocols, and your on-premises network is extended to Azure. The VPN connection can either be point-to-site or site-to-site.

A point-to-site VPN enables an individual computer or device to establish a connection to resources deployed in an Azure virtual network.

In a site-to-site VPN, a secure tunnel is established between your Azure virtual network VPN gateway and your own on-premises VPN device or router; this enables all devices from your on-premises network to communicate with resources in your Azure virtual network.

Azure ExpressRoute

Azure ExpressRoute offers a private network link between your on-premises network and Azure data centers. The traffic in Azure ExpressRoute does not traverse the public internet and hence offers improved security, reliability and lower latency compared to a VPN. It is a reliable hybrid connectivity service backed by Azure SLAs.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

For businesses that need reliable, high-throughput access to Azure for mission-critical workloads, ExpressRoute is the recommended solution.

Connectivity and routing

Azure creates default network routes to enable the communication of resources in its virtual network as well as to external resources. Each subnet will have its own associated routing table. The routing table will have entries for service/system endpoints as well as additional entries for Border Gateway Protocol (BGP) and routes from different connected virtual networks.

Azure offers the flexibility of creating user-defined routes (UDRs), also known as custom routes, to override the default routes. For example, you can create a UDR and route all traffic from a subnet to a network virtual appliance for further filtering or packet inspection.

If the network is communicating with an on-premises gateway through BGP, additional BGP routes are added in the routing table. Similar to traditional routers, the next hop of a packet is selected based on the longest prefix match. In the case of multiple entries that match the prefix, the routing happens in the order of UDRs, BGP, followed by system routes.

Load balancing

Load balancing services in Azure can be broadly classified as global/regional or HTTP(S)/non-HTTPS based services. The different load balancing services available in Azure are Azure Load Balancer, Application Gateway, Azure Traffic Manager and Azure Front Door.

Azure Load Balancer

Azure Load Balancer is used for layer four load balancing, for UDP and TCP protocols. It can be either internal or internet-facing. You can use TCP and HTTP-based rules to check the health of the backend pool and map the incoming connections to them.

Azure Application Gateway

Azure Application Gateway is a layer seven load balancer that operates at the application layer of the OSI model. This application delivery controller (ADC) as a service helps route traffic to and from your application based on the content of the HTTP requests. The routing can either be path-based or multi-site routing.

Azure Front Door

Azure Front Door is a load-balancing service best suited for multi-region application deployments, as it helps route traffic to the best available backend based on factors such as latency and availability. It offers features such as SSL offloading, caching, URL rewriting, etc. that help improve user experience.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Azure Traffic Manager

This DNS-based global load balancing service helps distribute application traffic across resources in multiple Azure regions. The routing method can be multi value or based on a defined priority, weight, performance, subnet and geography.

Content Delivery Service (CDN)

Though not a load balancing service, Azure CDN helps in accelerating application delivery by caching frequently accessed content across a network of strategically located servers known as edge nodes. It offers benefits such as bandwidth optimization, secure content delivery, dynamic site acceleration and improved application performance.

Application architecture

An effective application architecture should include optimal services to take care of messaging, event handling, caching, application configuration management and API integration. Azure has several out-of-the-box solutions that help achieve these requirements effectively.

Messaging solutions

Application components often communicate with each other through messages, which are nothing but raw data that need to be processed by the receiving component. The messaging services available in Azure are Azure Queue Storage and Azure Service Bus.

Azure Queue Storage

Azure Queue Storage is a scalable message service capable of holding millions of messages, which can be accessed using a REST interface. It is restricted only by the limits of the storage account where it is created.

This is a good application service when you need a simple queue service for messages. You can use it for tracking the progress of how a message is processed inside a queue and also to audit all the messages that are sent to the queue. You can use this service if the storage queue is expected to exceed 80 GB.

Azure Service Bus

Azure Service Bus, on the other hand, is an enterprise message broker for decoupling applications and services. Messages can be sent with Azure Service Bus using either message queues or publish-subscribe topics.

Message queues will have dedicated infrastructure and can hold messages until they are received by the target. In Azure Service Bus publish-subscribe topics, the message can be sent to multiple recipients, i.e., subscribers who have subscribed to a topic.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Event-driven architecture

Events are used for broadcast communications and are often lightweight when compared to messages. They use a publish-subscribe-based deployment model where the event is sent by the publisher and received by the subscribers. Azure Event Hubs is a managed messaging service that can receive and process millions of events at any given time.

Event Hubs is good for large-scale, real-time and batch data ingestion. It is helpful for applications that may have to handle large volumes of events such as with IoT applications or big data processing.

Event Hubs uses a pull model for messages; messages stay in Event Hubs and are not deleted after a read, thus remaining available for other consumers.

You can use Event Hubs for use cases such as live dashboarding, click streams, anomaly detection analysis, etc. The performance of the service depends on the chosen service tier, i.e., Basic, Standard or Premium; this is an important selection criterion and will depend on your application requirements.

Caching

Caching stores frequently accessed data to faster storage located closer to the application. Caching is for scenarios where you are not expecting many changes to the original data, meaning it will remain relatively static. It is also recommended to use caching when there are network latency concerns in terms of accessing the original content.

Azure Cache for Redis is a recommended service for caching requirements, as it offers an in-memory managed cache based on the popular Redis platform. You can either opt for the open-source implementation of the service, Redis Open Source (Redis OSS), or use the managed service from Redis Labs.

The service can be used as a distributed cache service, message broker or session store. It can also implement the cache-aside pattern, where only the required data from databases are loaded into the cache, or for cache static content from websites such as banners, headers and footers.

Application configuration management

Application configuration management separates the configuration of the application from code deployment and allows rapid adjustments to feature availability as needed. Separating off configuration as a service enables systems to dynamically manage the deployment lifecycle.

The Azure App Configuration service helps to centrally manage application settings, configurations and feature flags. It is delivered as a fully managed service, with native integration with popular development tools such as Visual Studio, Visual Studio Code, Azure CLI, etc. It comes pre-built with capabilities, such as feature flag management, resource tagging and configuration management with custom-defined dimensions.

The service also integrates with Azure AD for identity management and can be used in both development and production environments.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

API integration

With the increase in microservices and decoupled application architecture patterns, there has also been an increase in the number of APIs that get created or are integrated with applications. As the number of APIs increases, the complexity of managing and maintaining them also increases. Hence, it is important to incorporate a service that can act as a frontend for all the APIs used by an organization and help with centralized management, security and performance analytics.

Azure API Management is such a service. It helps to publish, secure, maintain and analyze all the APIs from a single pane of glass. You can control the usage of APIs with policies for usage quotas, call rate limits, request validation/transformation, etc.

It also helps implement compliance and security measures to prevent the malicious usage of APIs. As the number of APIs and the rate at which they are revised increases, it is imperative to implement Azure API Management to help deal with this complexity.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Design solutions using the Azure Well-Architected Framework

The right architecture is crucial for your cloud-based apps. It serves as the basis for all of your digital transformation efforts. The architecture's specifics change based on whether you want to use cloud-native services or plan to lift and shift an existing architecture to the cloud utilizing IaaS. Regardless of the strategy you select, a framework for architecture excellence is required to help you deploy your workloads correctly the first time.

The Azure Well-Architected Framework is a set of curated best practices, key principles and considerations structured under five pillars that help you design an Azure environment with confidence. The five pillars are cost optimization, operational excellence, performance efficiency, reliability and security.

These tenets help to ensure alignment of business and functional requirements with the right set of Azure services and best practices. They facilitate the proper balance of capabilities, cost, performance and risk while making Azure design decisions. This results in:

- Reduced wastage of cloud resources and improved cost savings.
- Effective application deployment and operational agility by leveraging DevOps practices.
- Applications designed to overcome performance bottlenecks, especially during peak usage time.
- Cloud architecture that is resilient to unplanned downtime as well as planned maintenance.
- Workloads that are protected against vulnerabilities and security gaps from the ground up.

The general design principles that the Azure Well-Architected Framework focuses on can be summarized as follows:

- Architecture should be flexible to evolve and adapt to new features introduced in services or better-suited tools and technologies.
- The architecture design should be data-driven, be it related to cost, performance or application usage.
- Organizations should promote a collaborative culture where different business units adopting cloud can educate and enable each other about best practices.
- Automation should be integrated into application development, deployment and day-to-day operations.

Let's look at the constructs of the Azure Well-Architected Framework in detail by exploring each of the five pillars introduced above.

Cost optimization

Cost optimization focuses on delivering the best return on investment in the cloud. As cloud services are delivered on demand, there is a paradigm shift from the traditional capital expense (CapEx) model that most organizations are used to.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

In the CapEx model, investments happen according to a defined timeline, after which they should be decommissioned. Cloud costs, on the other hand, are considered an operating expense (OpEx), i.e., an ongoing expenditure that depends on the usage of resources by the organization. This, in turn, has a direct impact on the balance sheet, leading to profit or loss.

With cloud computing, there is no need for a large upfront investment to host your workloads; however, the recurring monthly cost should be properly managed so that it remains within the organization's budget. To optimize cloud spend, you need to constantly monitor your cloud usage and take cost optimization measures that work well for your environment. This is an ongoing process, not a one-time fix.

The Azure Well-Architected Framework provides the following cost optimization best practices for organizations to manage their Azure costs.

Implement the right costing model

A cost estimate should only be done after you have finalized your services, including making sure you have the appropriate resource SKU sizes. You should also consider expected usage patterns, such as peak season usage, scaling, data download requirements and network traffic, that could add to costs.

The Azure pricing calculator then allows you to arrive at a monthly cloud bill estimate. This will act as a benchmark for your cloud spend and help you keep a close eye on any outliers during the billing cycle.

Optimize resources

With persistent usage patterns, evaluate cost-saving options such as bring your own license (BYOL), Azure Hybrid Benefit, reservations and spot virtual machines to bring down expenses. When it comes to service models, cloud-native and PaaS services are better for optimizing costs compared to IaaS. While a lift-and-shift migration to IaaS can be your short-term goal, plan on refactoring services to PaaS in the long run to reduce OpEx.

It is also important to select the right service tiers. You may opt to start with a lower tier of resources and scale up later if it doesn't suit your requirements, rather than starting with a higher tier and incurring unnecessary spend.

Use cost management tools

Azure offers cost management tools to drive visibility and accountability. For example, you can proactively set a budget and keep an eye on the actual and forecasted cost. You can also create alerts based on thresholds to analyze and implement optimizations if the cost is trending high against your determined budget.

You should also review the monthly billing statements to understand any resource usage anomalies. Tagging resources in Azure and reviewing the tags in each bill is a recommended approach to segregate costs, especially when there are multiple business units using the same Azure subscription.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Reduce wastage

Idle, unused and orphaned resources can inflate your cloud bills, as you end up paying for resources that are not being used productively. For example, you are required to pay for the disks associated with a virtual machine, even if it has been shut down. Also, storing infrequently used data in a storage hot tier will result in higher storage costs, while accessing large files over the network every time rather than caching them could spike network usage charges.

You can leverage the cost management feature in Azure Advisor to identify unused or idle resources and implement cost optimization recommendations.

Operational excellence

The Azure platform comes prebuilt with tools and capabilities that help provide end-to-end visibility of your application lifecycle and improve operational efficiency. By fine-tuning how you develop, deploy and manage your applications, you can improve the overall user experience.

Consider the following Well-Architected Framework guidelines to get the most out of your workloads deployed in Azure.

Modernize application delivery

You should integrate DevOps practices for continuous integration and delivery for efficient lifecycle management of your applications. It is not just the application you should automate and version, but the infrastructure deployment as well using infrastructure-as-code practices. Irrespective of whether you are deploying legacy applications or the latest microservices-based applications, you can adapt the principles of continuous integration and continuous deployment (CI/CD).

Focus on automating as much as possible to reduce human effort and error. More than a technical process, DevOps practices require a cultural change in organizations, primarily breaking down the silos within which development and operations teams often work.

Implement continuous monitoring

This is a common theme across different architecture pillars, as continuous monitoring is essential for fine-tuning existing processes and improving efficiency. You should instrument applications to collect metrics and logs and monitor them on a regular basis using Azure Monitor services.

Take advantage of the out-of-the-box default metrics and logs of Azure services to achieve a basic level of visibility. The monitoring approach should be multi-layered, where information is gathered from all associated components to gain visibility into application performance and health.

Fine-tune your test strategy

Define your application test strategy to proactively identify potential performance bottlenecks and security issues. Make sure to test inter-component communication as well as application dependencies. A robust test plan should cover end-to-end stack testing, starting from the infrastructure up to the application code.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Performance efficiency

Applications should be designed to meet peak usage demands without any performance bottlenecks. This requires scaling resources either vertically or horizontally, as well as fine-tuning applications to optimize performance. The following guidelines will help in achieving this.

Use auto scaling

Your application could hit performance issues due to resource constraints in peak usage periods. To avoid this, you have to make sure that the application is designed for scaling, wherein resources are allocated dynamically as usage increases. One way of doing this is via scale-out configurations, where additional instances are added to a service when usage exceeds defined thresholds.

Although not strictly autoscaling, you can consider vertical scaling as well, where the compute resources available to an instance are increased. However, this approach will be constrained by the maximum amount of CPU/memory that can be supported by the instance SKU.

Prioritize network and storage performance

Start by selecting the network and storage services that are best suited for your applications to avoid hitting performance snags. Consider aspects such as the location from which users access applications and the average network latency that can be tolerated by the application before selecting the Azure region. You can also place a messaging layer as a buffer if applications need additional time to process requests.

For the storage layer, consider partitioning to reduce contention and ensure scalability and consistent performance. Caching mechanisms also enable you to handle frequent data access requests, instead of overwhelming the storage layer and causing performance issues.

Optimize your applications

Bottlenecks in any layer of the application could impact overall performance. Hence, ensure that you have visibility into all layers through various channels such as infrastructure monitoring or application instrumentation. Resource utilization and the general health of all application components should be reviewed continuously to spot issues such as resource contention, dependency failures or faulty code. Implement continuous monitoring to identify the right patterns and optimization measures that work well for your application.

Reliability

Complex applications have several moving parts, and downtime in any of them could impact your entire application, leading to data loss and breach of SLAs. Possible causes of downtime include but are not limited to component failures, data center outages and partial network outages.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Ensure high availability

Application architecture should be designed for high availability to avoid downtime caused by outages. The basic design principle is to eliminate single points of failure across all components and dependencies. Systems should be resilient to failures and recover automatically from them wherever possible.

Start by defining the SLA required for the application, and make sure you have redundancy built into all components of the application to meet this SLA. This can be achieved through design patterns such as clustering, automatic failovers and load balancing.

Strategize for failure recovery

Designing for reliability should also include options for recovering from data loss and large-scale outages. Start by defining your RPO and RTO goals and implement recovery strategies aligned with these. You might be required to evaluate multiple options before you can zero in on the best strategy for your business.

A cost/benefit analysis is also a must before deciding on a solution. Make sure to factor in automation wherever possible to bring down the RTO. Recovering from failures requires a holistic approach that includes a backup/restore as well as disaster recovery solution.

You can read more about the different approaches to ensure high availability in the "Design business continuity solutions" section of this eBook.

Security

Ensuring the security of your application and data in the cloud is of utmost importance, as any breach of your cloud environment could lead to long-lasting impacts on your business. At the end of the day, it is about making sure that the organization's assets do not fall into the wrong hands. A security breach, in addition to causing financial loss, can also damage the reputation of your business, impacting future prospects.

The rule of thumb is not to consider security as an afterthought, but instead bake it into all components of your applications and all stages of your application lifecycle.

Follow the shared security model

The basic premise of security in Azure, as with any other cloud service provider, is the shared responsibility model. While Azure provides the required security configurations for all services, it is up to the customer to ensure that they fine-tune it per their specific needs. This will vary depending on the cloud model being selected.

For example, in a PaaS deployment, the security of the OS, hosts and infrastructure is handled by Microsoft, but it is up to the customer to implement guardrails around IAM, data governance and client endpoint security. The shared responsibility model in the context of various deployment models is shown below:

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Implement defense-in-depth

Azure recommends a layered approach for security, where all components starting from the infrastructure are configured to improve security posture. This is also commonly known as defense-in-depth security.

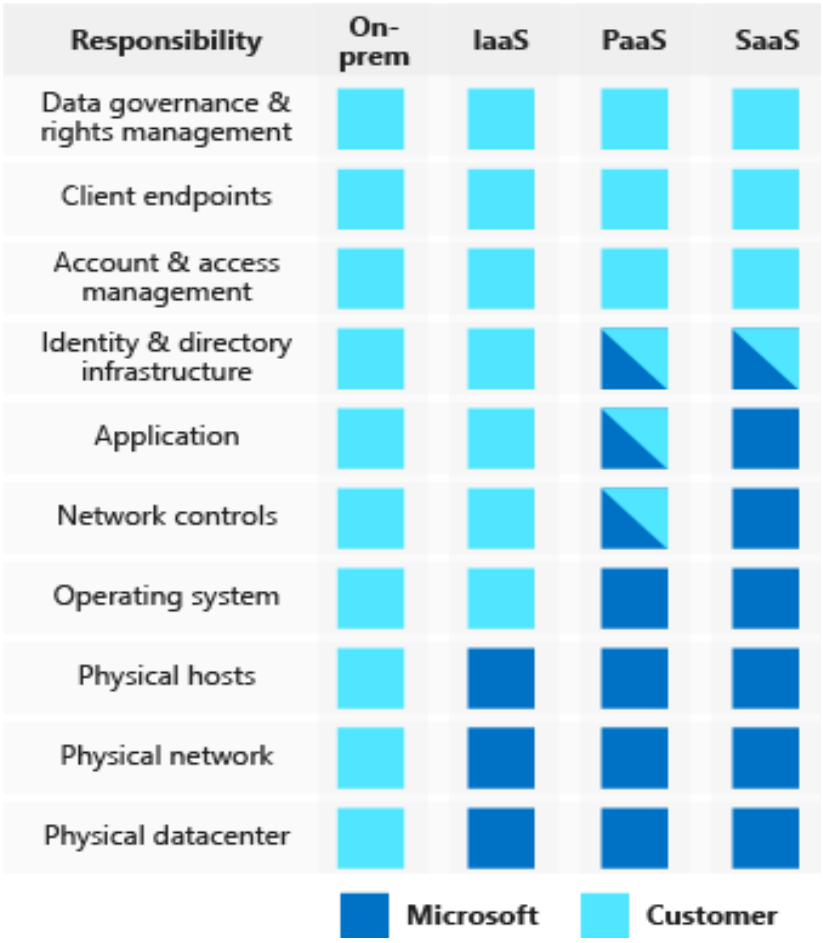


Figure 1: Shared responsibility model in Azure (Source: Microsoft)

With security measures implemented at all layers, you will have better protection against the lateral movement of threats.

You also need to keep in mind that no system can be made 100% secure, as the threat landscape in the cloud keeps evolving. Hence, you should continuously monitor your Azure environments to identify threat patterns and infiltration attempts and update your security configurations.

Also, make sure to safeguard your data through encryption. This can either be Azure-managed encryption or customer-managed encryption with keys protected in Azure KeyVault.

Below are some of the common attack patterns that you should be aware of in order to plan your security strategy.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Compute layer exploits

The application layer can be attacked by exploiting vulnerabilities in the code or compromised binaries. Common attack patterns are cross-site scripting, SQL injection and compromised binary exploitation. You can leverage Azure Web Application Firewall (WAF) or advanced network virtual appliances to handle these attacks.

VMs or compute layers can also face security breaches including but not limited to malware attacks, malicious code injection and OS vulnerability exploitation. You can leverage built-in tools such as Azure Defender, Microsoft Antimalware and Azure Automation Update Management to mitigate such threats.

Distributed denial-of-service (DDoS) attacks

DDoS is a common form of security threat where attackers continuously flood target systems with requests to overwhelm the network resources, rendering the application unresponsive. Azure offers comprehensive DDoS protection resources you can leverage for comprehensive perimeter protection.

Identity and credential theft

This is another common form of attack used by malicious actors to gain access to your system. Credential theft in systems can also lead to the lateral movement of a threat, which could then go undetected. Hackers can take advantage of a vulnerability in your authentication protocol or use methods such as social engineering or phishing to get ahold of admin credentials to enter the environment.

In addition to following the principle of least privilege with RBAC, you should also implement measures such as context-aware conditional access and just-in-time access to protect your workloads.

Focusing on the above five architecture pillars will help in building a solid foundation while deploying or migrating your workloads to the cloud. They cover all aspects of an architecture for robust, scalable and secure applications that can be trusted by customers.

On top of following the Azure Well-Architected Framework, designing a secure Azure cloud environment requires a mindset and culture change. It is important to keep remember that data protection is everyone's job. Ensuring the security of your mission-critical data — be it from data theft or corruption/loss — should be of utmost importance while designing Azure environments.

Cloud and data protection teams have a direct stake in protecting mission-critical applications. You can use the built-in solutions offered by Azure or market-leading third-party solutions to achieve this. As an Azure architect with AZ-305 certification, you will be expected to recommend the best solution to ensure comprehensive data protection.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

Conclusion

With cloud computing evolving rapidly, it is crucial to expand your skills and keep them up to date so that you are always ahead of the curve.

The AZ-305 evaluates and certifies your skills in designing and implementing robust solutions using the Microsoft Azure platform. Along with Azure documentation, you can leverage this guide to achieve a good understanding of the key focus areas of the exam.

The topics covered here should be considered as a baseline to prepare for the exam. Candidates should have a solid grasp of the concepts covered here, especially those around data security and protection, to pass the AZ 305 exam. However, in addition to this guide, make sure to go through Azure's own documentation for each of the modules discussed to further familiarize yourself with the services.

It is also recommended to get hands-on practice in an Azure environment. You can either create an Azure free account or use the Azure sandbox environments in Microsoft Learn. There are instructor-led training courses as well you can take advantage of to prepare for the exam.

Azure also offers sample questions to help you measure your knowledge and boost your confidence to face the actual exam questions.

All the very best for your exam!!

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.

About the Author



Leah Troscianecki

Leah is a Principal Product Marketing Manager for Veeam's Public Cloud solutions, focused on driving adoption of better together solutions for backup to and backup within the AWS, Microsoft Azure and Google Cloud ecosystem of users. Her specializations include resale, managed service and OEM partnerships, business continuity and disaster recovery (BC/DR) solutions, and IaaS/PaaS/SaaS offerings. On the weekends, Leah can be spotted cheering on the Ohio State University Buckeyes at football games in Columbus, Ohio.

To learn about Veeam's Azure backup and recovery solutions, visit this webpage: https://www.veeam.com/backup-azure.html.

© 2023 Veeam Software. Confidential information. All rights reserved. All trademarks are the property of their respective owners.