

Azure Onboarding Guide for IT Organizations

Azure Onboarding Guide for IT Organizations

7-Jul-16

Version 1.3

Prepared by

Joachim Hafner

Authors and Contributors

The following resources contributed to this version of the Azure Onboarding Guide:

Author

Joachim Hafner – Cloud Solution Architect at Microsoft

Contributors and Reviewers

Carsten Lemm – Cloud Solution Architect at Microsoft

Eduardo Kassner – Director of Cloud Solution Architecture at Microsoft

Barry Briggs – Independent Consultant for software end enterprise computing

Cloud Solution Architects Microsoft Germany

© 2016 Microsoft. All rights reserved. This document is for informational purposes only. Microsoft makes no warranties, express or implied, with respect to the information presented here.

Table of Contents

1	Introduction	5
2	Moving to the cloud	6
2.1	Adaptation of the IT organization	7
2.2	Transforming the IT organization	7
2.3	Adopting the cloud	9
2.4	Preparing and training IT staff for the cloud	13
2.5	Recommendations for moving to the cloud	15
3	Managing security, compliance and data privacy	17
3.1	Working to keep customer data safe	18
3.1.1	Security design and operations	18
3.1.2	Infrastructure protection	20
3.1.3	Network protection	21
3.1.4	Data protection	22
3.1.5	Identity and access	23
3.2	Owning and controlling data	24
3.3	Managing compliance and data privacy regulations	25
3.4	Azure Security Center	26
3.5	Microsoft US government cloud	27
3.6	Microsoft cloud in Germany	27
3.7	Cloud security recommendations for enterprise architects	28
4	Azure enterprise administration	36
4.1	Understanding Azure subscriptions	37
4.2	Managing Azure subscriptions	39
4.3	Defining naming conventions	43
4.4	Recommendations for Azure enterprise administration	43

5	Integrating Azure into the corporate network.....	45
5.1	Choosing the right connectivity option	45
5.1.1	Using ExpressRoute.....	48
5.1.2	Using Site-to-Site VPN.....	52
5.2	Protecting virtual networks.....	53
5.2.1	Network Security Groups	53
5.2.2	Forced tunneling	54
5.2.3	Virtual Appliances	54
5.3	Routing of network traffic	56
5.4	Managing public and private IP addresses.....	57
5.5	Recommendations for cloud connectivity	58
6	Extending Active Directory to Azure.....	60
6.1	Synchronizing/federating Active Directory Domain Services with Azure AD	61
6.2	Working with multiple forests and domains.....	62
6.3	Multi-Factor Authentication	65
6.4	Hosting Active Directory domain services	66
6.5	Using additional Azure Active Directory elements	67
6.5.1	Azure AD B2B Collaboration	67
6.5.2	Azure AD B2C Collaboration.....	68
6.5.3	Azure AD Domain Services.....	68
6.5.4	Azure Application Proxy	69
6.6	Recommendations for using Azure Active Directory.....	69
7	Operating Azure IaaS Services	71
7.1	Gaining operational insights	71
7.1.1	Getting started with Log Analytics	71
7.1.2	Creating log searches and raising of alerts	74
7.1.3	Securing data.....	76
7.2	Backing up and restoring data	76

7.2.1	Azure virtual machines.....	77
7.2.2	Files and folders.....	80
7.2.3	Enterprise applications.....	82
7.3	Establishing secure remote access.....	84
7.4	Automating operational procedures.....	87
7.5	Managing IT services according to ITIL.....	88
7.6	Recommendations for operating Azure IaaS Services.....	89
8	Migrating existing services to Azure.....	90
8.1	Configuring virtual machine and application migrations.....	90
8.2	Mapping of networks and subnets	91
8.3	Planning and testing failover	92
8.4	Recommendations for migrating existing services to Azure	92
9	Offering management for cloud-based services.....	93
9.1	Consuming services	94
9.2	Provisioning of cloud services.....	95
9.3	Metering consumption per application	97
9.4	Billing and price prediction	98
9.5	Managing the lifecycle.....	99
9.6	Recommendations for cloud service provisioning.....	100

1 Introduction

There are a lot of good reasons for enterprises to move to the cloud, such as greater business agility, keeping track with the speed of innovation, and cost savings. The current state of the various cloud surveys shows that cloud adoption is growing and has now hit its stride. The strong growth in the use of cloud means that the majority of organizations are now operating in a hybrid environment that consists of on-premises and cloud-based services.

The cloud is also changing how companies consume technology. Employees and business departments are more empowered than ever before to find and use cloud applications, often with limited or no involvement from the IT department, creating what's called "shadow IT."

Despite the benefits of cloud computing, companies face numerous challenges including the integration of cloud services into the enterprise architecture, security and compliance of corporate data, managing employee-led cloud usage, establishing operational processes for cloud services, and even the development of necessary skills needed in the cloud era.

As companies move data to the cloud, IT departments are looking to put in place policies and processes so that employees and business departments can take advantage of cloud services that drive business growth without compromising the security, compliance, and governance of corporate data.

The purpose of this document is to provide an overview, guidance, and best practices for enterprise IT departments to introduce, consume, and manage Microsoft Azure-based services within their organization. The target audience is enterprise architects, cloud architects, system architects, and IT managers.

This document is not intended to replace existing documentation about Microsoft Azure services and features.

2 Moving to the cloud

The evolution and maturation of cloud technologies have brought enterprise IT into a transitional stage. A 2015 [EDUCAUSE](#) study found that CIOs expect a significant shift in focus in the next five years, away from managing primarily infrastructure and toward the cloud.

But why is cloud technology so irresistible? Why is making the migration such a good idea for businesses? Answering this question for your business before you make the move is essential. There are no two ways about it: Your business will move to the cloud, and making that move is a good idea. But your success hinges on your reasoning for making the move. There are many different opportunities to migrate to the cloud, each of which may have a different reason behind it, and it's critical that you identify each of these reasons.

According to a study from Accenture ([Behind Every Cloud, There's a Reason](#)), there are six most common business and technology drivers for making the move to the cloud, including how to identify these drivers, how to identify the right drivers for your program, and how to define your drivers. Properly identifying, defining, and balancing these drivers can help your business successfully execute its cloud strategy and move toward a true transformation.

The six drivers for making the move to the cloud fall under two main categories: business drivers, including business growth, efficiency, and experience, and technology drivers, including agility, cost, and assurance.

Business growth

What are you doing to make your company more successful from the perspective of expansion? This can take a number of different forms (for example, driving sales, enlarging wallet share, or increasing productivity), so it's important to clearly define how you intend to achieve this growth.

Efficiency

Efficiency includes streamlining processes to get work done faster or with less resources. This can turn around and fuel growth (for example, by allowing you take on more work) or reduce costs (for example, by reducing the amount of resources required).

Experience

Whether it's external or internal, the customer experience is of utmost importance in today's world. A good experience can increase brand loyalty among customers and retention among employees. In general, a positive customer experience is strongly tied to brand value.

Agility

Agility is the most common cloud driver today, especially when IT is leading the charge. Being agile helps IT change and scale fast enough to keep up with business needs.

Cost

The difference between the technology driver cost and the business driver efficiency is often misunderstood. Although efficiency can lead to cost savings, the cost driver focuses on reducing the cost of IT licenses or operations and/or redefining the cost model for technology solutions.

Assurance

Finally, assurance encompasses the achievement of mission-critical technology outcomes, such as protecting against datacenter crashes or security breaches and maximizing disaster recovery effectiveness. Going to the cloud for assurance frees IT to be more strategic and passes these responsibilities to a provider that is typically better at handling them than your business is.

As you begin your move to the cloud, it's important to identify which of these factors is driving your journey. Doing so should help inform your next steps and justify making the move.

2.1 Adaptation of the IT organization

The effective adoption of cloud services requires changes to an organization's existing operational practices and procedures (see EDUCAUSE [Preparing the IT Organization for the Cloud](#)). The external nature of cloud services may require an organization to rethink its IT service management and disaster recovery practices, as well as how given cloud services integrate with its existing in-house technology infrastructure. The pay-as-you-go cost model common with cloud services may entail changes to financial management practices and total cost of ownership calculations. Procurement processes may need to be adjusted to increase agility and effectively address the unique risks associated with cloud service, and new vendor management roles may need to be established and resourced to ensure ongoing compliance with contract terms.

2.2 Transforming the IT organization

Cloud strategy development is an evolutionary process in most enterprises. Adopting a cloud strategy requires careful coordination among a variety of stakeholders, including IT and information security staff, legal teams, compliance experts, procurement specialists, and institutional leadership. Once an enterprise cloud strategy is adopted, the implementation of those strategies requires transformation in the IT organization. Some common approaches and stages to developing an enterprise-wide cloud strategy include:

- **Cloud aware**

Enterprise users and IT staff are aware of broad cloud trends but are not yet prepared to adopt public-cloud solutions. These institutions may choose to build on-premises solutions in a way that prepares them for an eventual move to the cloud.

- **Cloud experimentation**

The IT organization begins to learn about the various cloud services available to them in the forms of SaaS, PaaS, and IaaS. The organization may begin deploying some common SaaS solutions (such as Office 365), which sometimes grows into testing IaaS deployments.

- **Opportunistic cloud**

The IT organization begins to actively seek out cloud solutions that meet new business requirements. Services may remain as traditional on-premises deployments, but cloud solutions are considered and deployed when reliability, scalability, or other benefits are perceived.

- **Cloud first**

This strategy places cloud at the top of the decision-making chain. The default assumption within the enterprise is that cloud services will fulfil the majority of the enterprise computing needs.

The adoption of the various cloud strategies causes a paradigm shift that impacts both the IT organization and IT staff members. Business units are increasingly driving the selection and adoption of cloud IT solutions, and in doing so they may bypass the IT units. Enterprises will achieve the best outcomes when their IT organizations serve as enablers, simplifying and accelerating business units' adoption of cloud services. In order to evolve into the role of cloud enabler, IT units should carefully consider the value they bring with regard to cloud service adoption, such as:

- Establish strategy and goals
- Define criteria for moving to or starting applications in the cloud
- Architect core infrastructure components for cloud integration: Identity, Networking, Security
- Acquire cloud development skills
- Retool for adoption and change management
- Take a systematic and disciplined approach to security and compliance

Now more than ever, IT units must clearly understand the evolving needs of business and units and be prepared to help them assess the full range of possible solutions to their technology needs. Successful IT organizations will find ways to simplify and accelerate cloud adoption by reducing the barriers their partners face and by helping their company avoid potential pitfalls.

IT organizations must develop competencies with cloud technologies and services even as those services evolve and change. Practically, this means that staff must have time to explore new

technologies and that organizations may need to increase their investment in IT staff training. IT organizations that fail to provide sufficient time for training and exploration will likely find themselves unable to contribute meaningfully to the campus technology conversation. Business units will not wait. They will simply bypass IT organizations unable to meet their needs. Change management practices and IT governance processes need to become agile and rapidly responsive to the needs of users, while still assessing risks to the organization.

2.3 Adopting the cloud

In any transformative change, it is important to understand what the destination is and what the waypoints along the journey will be. There are multiple potential destinations for any application, and IT cloud deployments will be a mixture of them (hybrid cloud). Hybrid cloud uses compute or storage resources on your on-premises network and in the cloud. You can use hybrid cloud as a path to migrate your business and its IT needs to the cloud or integrate cloud platforms and services with your existing on-premises infrastructure as part of your overall IT strategy.

- **Private Cloud**

Private Cloud technologies are hosted in an on-premises datacenter or in a datacenter of a managed service provider. This might be necessary for certain applications or data that can't be moved to the cloud. Private Clouds are especially useful if they implement a technology stack that is consistent with the Public Cloud. Microsoft Azure Stack is a product that enables organizations to deliver Azure services from their own datacenter. It helps you build and deploy your applications the same way regardless of whether it runs on Azure or Azure Stack.

- **Infrastructure as a Service (IaaS)**

Infrastructure as a Service abstracts hardware (server, storage, and network infrastructure) into a pool of computing, storage, and connectivity capabilities that are delivered as services for a usage-based (metered) cost. IaaS services allow you to build and run server-based IT workloads in the cloud, rather than in your on-premises datacenter. IaaS services typically consist of an IT workload that runs on virtual machines that is transparently connected to your on-premises network. Your on-premises users will not notice the difference. IaaS is one of the most common cloud deployment patterns to date. It eliminates the need for capital expense budgets and reduces the time between purchasing and deployment to almost nothing. In addition, it is most similar to how IT operates today.

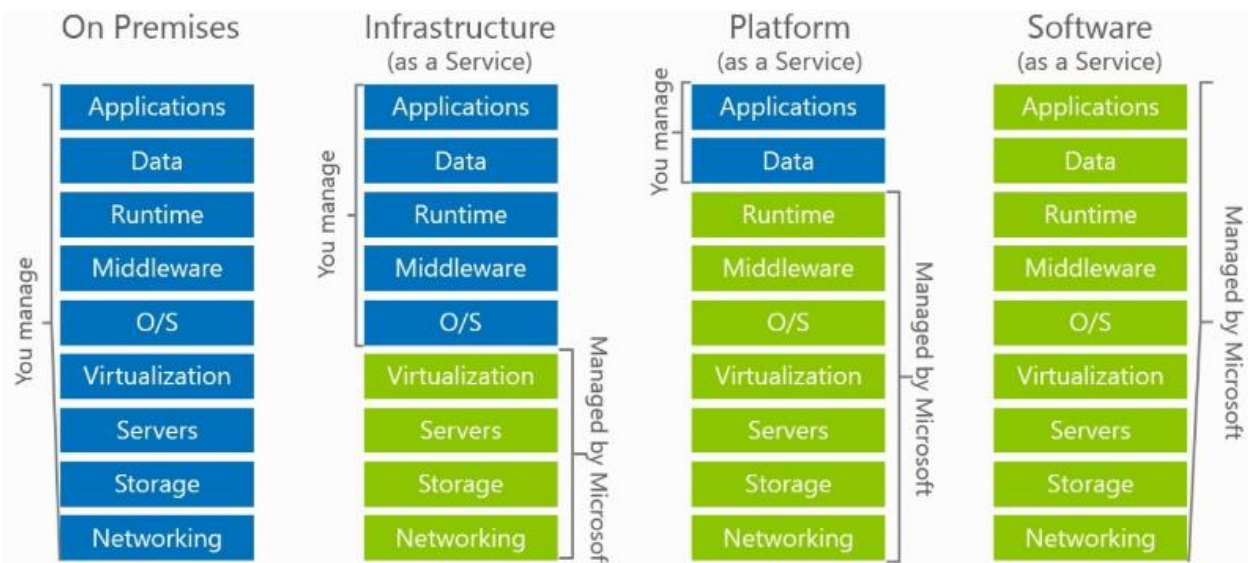
- **Platform as a Service (PaaS)**

Platform as a Service delivers application execution services, such as application runtime, storage, and integration, for applications written for a prespecified development framework. In a PaaS deployment model, enterprises are focusing on deploying their application code into PaaS services. PaaS provides an efficient and agile approach to operate scale-out applications in a predictable and cost-effective manner. Service levels and operational risks are shared because the consumer must take responsibility for the stability, architectural compliance, and overall operations of the application while the provider delivers the platform capability (including the infrastructure and operational functions) at a predictable service level and cost.

- **Software as a Service (SaaS)**

Software as a Service delivers business processes and applications, such as CRM, collaboration, and email, as standardized capabilities for a usage-based cost at an agreed, business-relevant service level. SaaS provides significant efficiencies in cost and delivery in exchange for minimal customization and represents a shift of operational risks from the consumer to the provider. All infrastructure and IT operational functions are abstracted away from the consumer. IT departments need only to take care of provisioning users and data and perhaps integrating the application with Single Sign-On.

The chart below shows the different responsibilities for IaaS, PaaS, and SaaS.



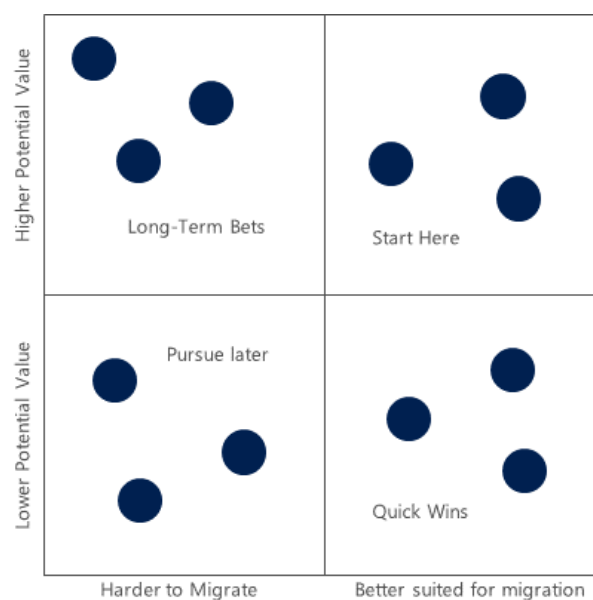
Most enterprises developed already or started developing new modern cloud applications. Those applications have been designed for the cloud right from the beginning, and they make use of PaaS offerings such as Azure Web Apps, Mobile Apps, Logic Apps, SQL Databases, Stream

Analytics, and HD Insight, among others. But there is also a huge amount of traditional enterprise IT that can benefit from the cloud as well without re-architecting existing applications.

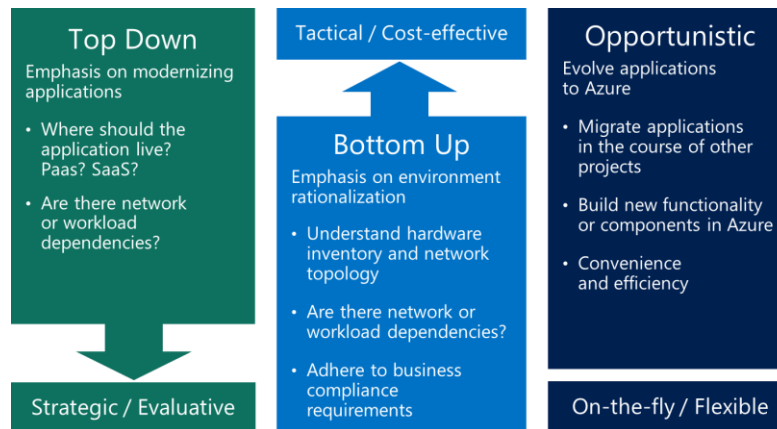
Large enterprises are running hundreds or thousands of applications running perhaps on tens of thousands of virtual machines. Key questions that need to be answered are: Which applications could be moved? How could they be moved? How to prioritize? How does it affect the business?

Therefore, it is important to create a well-attributed catalog of applications managed by IT. The relative importance of attributes such as business criticality, amount of integrations points, performance requirements, etc., can be weighted and a prioritized list can be built. You can think about those characteristics top-down or bottom-up. Top-down describes where each application should go to; bottom-up describes where it can go.

The top assessment first evaluates the security and compliance aspects. Then it assesses the complexity, interfaces, authentication, data structure, latency requirements, and other aspects of the application architecture. Next are operational requirements, service levels, maintenance windows, monitoring, and others. The result is a score that reflects the relative difficulty to migrate the applications. Furthermore, the financial benefits of the application such as operational efficiency, total cost of ownership, return on investment, and others have to be taken into account. The seasonality, required scalability, and elasticity of the application need to be considered and finally business continuity and resilience requirements that the application might have. With this analysis you can figure out the applications that have the highest potential value and are better suited for migrations.



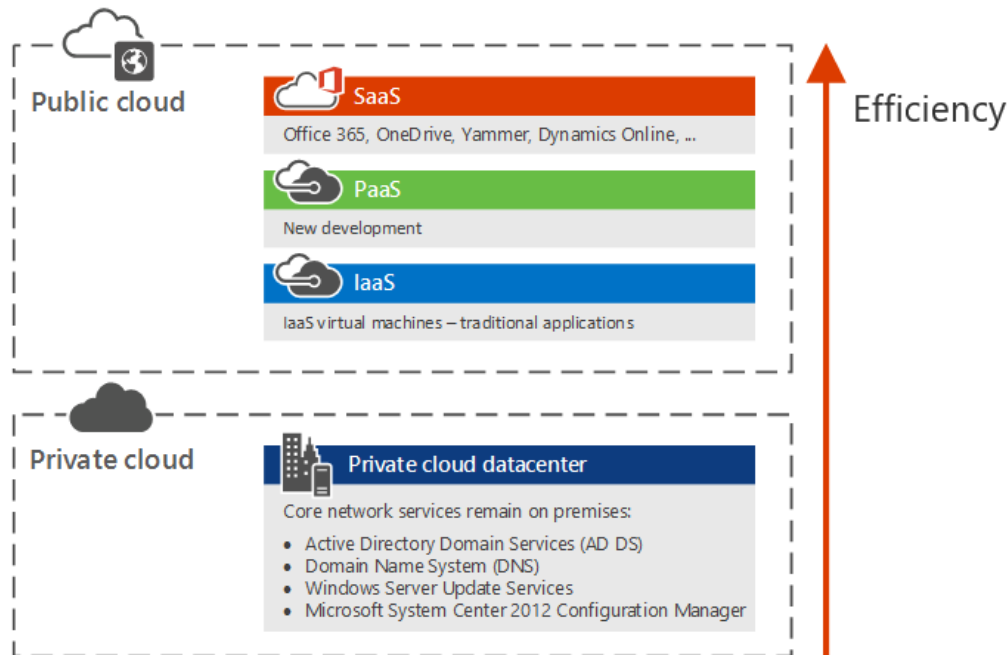
Analyzing the applications from a bottom-up perspective is aimed at providing a view into the eligibility, at a technical level of an application to migrate. Evaluated requirements are typically maximum memory, number of cores, operating system and data disk space, network interface cards, network and IP settings, load balancing, clustering, versions of operating systems, databases, middleware products, and web servers, among others.



Another aspect is the cloud platform, IaaS, PaaS, SaaS that the application should be migrated to. Many enterprise organizations take a three-step approach to cloud adoption. The first priority is to take advantage of SaaS productive workloads such as Office 365. The second priority is to base new modern cloud applications on PaaS (Azure SQL databases, Azure Web Apps, Logic Apps, Mobile Apps, etc.). The focus is on functionality rather than infrastructure. The third priority is moving existing applications to IaaS virtual machines by using one of the two approaches:

- **Lift and Shift:** Existing virtual machines are shifted to the cloud.
- **Build in the cloud:** applications are prebuilt in Azure, and traditional methods are used to back up and restore data.

To maximize efficiency, organizations are intending to use the higher level services of Azure wherever possible. Even though migration to Azure is the goal, retaining core network services in traditional on-premises datacenters will be necessary for the near future and results in a Hybrid Cloud.



This guide is focusing on getting Azure ready to use for PaaS and IaaS and provides best practices how to migrate, manage, and operate IaaS-based workloads in Azure. For further details about migration planning, please refer to this free ebook:

<https://info.microsoft.com/enterprise-cloud-strategy-ebook.html>.

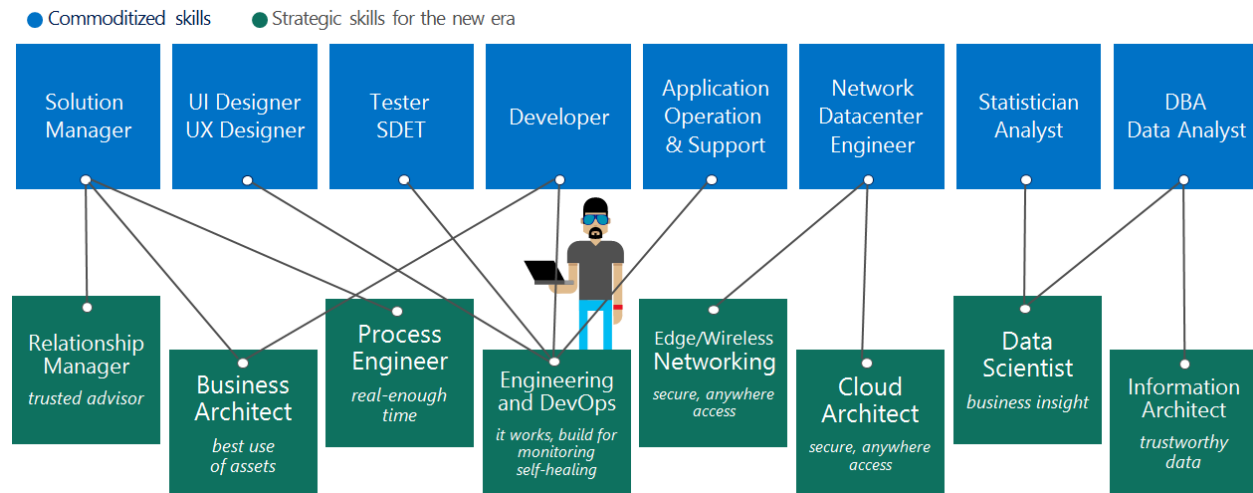
2.4 Preparing and training IT staff for the cloud

In order for IT staff to function as change agents supporting current and emerging cloud technologies, their buy-in for the use and integration of these technologies is needed. For this, staff need three things:

- An understanding of their roles and of any changes to their current position
- Time and resources to explore the technologies
- An understanding of the business case for the technologies

At each evolutionary phase during the history of the IT industry, the most notable industry changes are often marked by the changes to staff roles. During the transition from mainframes to the client/server model, the role of the computer operator largely disappeared, replaced by the system administrator. When the age of virtualization arrived, the requirement for individuals working with physical servers diminished, replaced with a need for virtualization specialists. Similarly, as institutions shift to cloud computing, roles will likely change again. For example, datacenter specialists might be replaced with cloud financial analysts. Even in cases where IT job titles have not changed, the daily work roles have evolved significantly.

IT staff members may feel anxious about their roles and positions as they realize that a different set of skills is needed for the support of cloud solutions. But agile employees who explore and learn new cloud technologies don't need to have that fear. They can lead the adoption of cloud services and help the organization understand and embrace the associated changes. Typical mappings of roles are:



Microsoft and partners offer a variety of options for all audiences to develop their skills with Microsoft Azure services.

- Microsoft Virtual Academy (<https://mva.microsoft.com/product-training/microsoft-azure>) offers training from the people who helped to build Microsoft Azure. From the basic overview to deep technical training, IT staff will learn how to leverage Microsoft Azure for their business.
- Microsoft IT Pro Cloud Essentials (<https://www.itprocloudessentials.com>) is a free annual subscription that includes cloud services, education, and support benefits. IT Pro Cloud Essentials provides IT implementers with hands-on experience, targeted educational opportunities, and access to experts in areas that matter most to increase knowledge and create a path to career advancement.
- The Microsoft IT Pro Career Center (<https://www.itprocareercenter.com>) is a free online resource to help map your cloud career path. Learn what industry experts suggest for your cloud role and the skills to get you there. Follow a learning curriculum at your own pace to build the skills you need most to stay relevant.

We recommend turning knowledge of Microsoft Azure into official recognition with Microsoft Azure certification training and exams. (<https://www.microsoft.com/en-us/learning/mcsd-azure-architect-certification.aspx>).

2.5 Recommendations for moving to the cloud

Identify your drivers to move to the cloud and take a systematic approach.

Recommendations for moving to the cloud	
Catalog existing applications	To understand what applications should be moved, when and how, it's important to create a well-attributed catalog of applications managed by IT. Then, the relative importance of each attribute (say, business criticality or amount of system integration) can be weighted and the prioritized list can be built.
Define criteria for moving to or starting applications in the cloud	You should set priorities within your migration plan based on a combination of business factors, hardware/software factors, and other technical factors. Your business liaison team should work with the operations team and the business units involved to help establish a priority listing that is widely agreed upon. For sequencing the migration of your workloads, you should begin with less-complex projects and gradually increase the complexity after the less-complex projects have been migrated.
Architect core infrastructure components for cloud integration	<p>You must account for the following elements when planning and implementing hybrid cloud scenarios.</p> <p>Networking for hybrid cloud scenarios includes the connectivity to Microsoft cloud platforms and services and enough bandwidth to be performant under peak loads.</p> <p>Identity for SaaS and Azure PaaS hybrid scenarios can include Azure AD as a common identity provider, which can be synchronized with your on-premises Windows Server AD, or federated with Windows Server AD or other identity providers. You can also extend your on-premises Identity infrastructure to Azure IaaS.</p> <p>Security for hybrid cloud scenarios includes protection and management for your identities, data protection, administrative privilege management, threat awareness, and the implementation of governance and security policies.</p>
Acquire cloud development skills	You must develop competencies with cloud technologies and services even as those services evolve and change. Practically, this means that staff must have time to explore new technologies and that you may need to increase your investment in IT staff training.

Retool for adoption and change management	Rethink your IT service management and disaster recovery practices, as well as how a given cloud service integrates with your existing in-house technology infrastructure. Consider the usage of cloud-based IT service management solutions.
Take a systematic and disciplined approach to Security, Governance, Compliance	<p>Invest in core capabilities within your organization that lead to secure environments:</p> <ul style="list-style-type: none">• Governance & Security Policy• Administrative Privilege Management• Identity Systems and Identity Management• Threat Awareness• Data Protection

3 Managing security, compliance and data privacy

Every business has different needs and every business will reap distinct benefits from cloud solutions. Still, customers of all kinds have the same basic concerns about moving to the cloud. They want to retain control of their data, and they want that data to be kept secure and private, all while maintaining transparency and compliance. What customers want from cloud providers is:

- **Secure our data**

While acknowledging that the cloud can provide increased data security and administrative control, IT leaders are still concerned that migrating to the cloud will leave them more vulnerable to hackers than their current in-house solutions.

- **Keep our data private**

Cloud services raise unique privacy challenges for businesses. As companies look to the cloud to save on infrastructure costs and improve their flexibility, they also worry about losing control of where their data is stored, who is accessing it, and how it gets used.

- **Give us control**

Even as they take advantage of the cloud to deploy more innovative solutions, companies are very concerned about losing control of their data. The recent disclosures of government agencies accessing customer data, through both legal and extra-legal means, make some CIOs wary of storing their data in the cloud.

- **Promote transparency**

While security, privacy, and control are important to business decision-makers, they also want the ability to independently verify how their data is being stored, accessed, and secured.

- **Maintain compliance**

As companies expand their use of cloud technologies, the complexity and scope of standards and regulations continue to evolve. Companies need to know that their compliance standards will be met, and that compliance will evolve as regulations change over time.

3.1 Working to keep customer data safe

3.1.1 Security design and operations

Secure cloud solutions are the result of comprehensive planning, innovative design, and efficient operations. Microsoft makes security a priority at every step, from code development to incident response.

- **Design for security from the ground up**

Azure code development adheres to Microsoft's Security Development Lifecycle (SDL). The SDL is a software development process that helps developers build more secure software and addresses security compliance requirements while reducing development cost. The SDL became central to Microsoft's development practices a decade ago and is shared freely with the industry and customers. It embeds security requirements into systems and software through the planning, design, development, and deployment phases.

- **Enhancing operational security**

Azure adheres to a rigorous set of security controls that governs operations and support. Microsoft deploys combinations of preventive, defensive, and reactive controls including the following mechanisms to help protect against unauthorized developer and/or administrative activity:

- Tight access controls on sensitive data, including a requirement for two-factor smartcard-based authentication to perform sensitive operations.
- Combinations of controls that enhance independent detection of malicious activity.
- Multiple levels of monitoring, logging, and reporting.


In addition, Microsoft conducts background verification checks of certain operations personnel and limits access to applications, systems, and network infrastructure in proportion to the level of background verification.

- **Assume breach**

One key operational best practice that Microsoft uses to harden its cloud services is known as the "assume breach" strategy. A dedicated "red team" of software security experts simulates real-world attacks at the network, platform, and application layers, testing Azure's ability to detect, protect against, and recover from breaches. By constantly challenging the security capabilities of the service, Microsoft can stay ahead of emerging threats.

- **Incident management and response**

Microsoft follows a five-step incident response process when managing both security and availability incidents for the Azure services. The goal for both types is to restore normal service security and operations as quickly as possible after an issue is detected and an investigation is started. The response is implemented using a five-stage process illustrated in the figure below, which shows the following activities: Detect, Assess, Diagnose, Stabilize, and Close. The Security Incident Response Team may move back and forth between diagnose and stabilize as the investigation progresses.



Detect

First indication of an event investigation

Assess

An on-call incident response team member assesses the impact and severity of the event. Based on evidence, the assessment may or may not result in further escalation to the security response team.

Diagnose

Security response experts conduct the technical or forensic investigation, and identify containment, mitigation, and workaround strategies. If the security team believes that customer data may have become exposed to an unlawful or unauthorized individual, parallel execution of the Customer Incident Notification process begins in parallel.

Stabilize, recover

The incident response team creates a recovery plan to mitigate the issue. Crisis containment steps such as quarantining impacted systems may occur immediately and in parallel with diagnosis. Longer term mitigations may be planned, which occur after the immediate risk has passed.

Close/post-mortem

The incident response team creates a post-mortem that outlines the details of the incident, with the intention to revise policies, procedures, and processes to prevent a reoccurrence of the event.

3.1.2 Infrastructure protection

Azure infrastructure includes hardware, software, networks, administrative and operations staff, and the physical datacenters that house it all. Azure addresses security risks across its infrastructure.

- **Physical security**

Azure runs in geographically distributed Microsoft facilities, sharing space and utilities with other Microsoft Online Services. Each facility is designed to run 24x7x365 and employs various measures to help protect operations from power failure, physical intrusion, and network outages. These datacenters comply with industry standards (such as ISO 27001) for physical security and availability. They are managed, monitored, and administered by Microsoft operations personnel.

- **Monitoring and logging**

Centralized monitoring, correlation, and analysis systems manage the large amount of information generated by devices within the Azure environment, providing continuous visibility and timely alerts to the teams that manage the service. Additional monitoring, logging, and reporting capabilities provide visibility to customers.

- **Update management**

Security update management helps protect systems from known vulnerabilities. Azure uses integrated deployment systems to manage the distribution and installation of security updates for Microsoft software. Azure uses a combination of Microsoft and third-party scanning tools to run OS, web application, and database scans of the Azure environment.

- **Antivirus and antimalware**

Azure software components must go through a virus scan prior to deployment. Code is not moved to production without a clean and successful virus scan. In addition, Microsoft provides native antimalware on all Azure VMs. Microsoft recommends that customers run some form of antimalware or antivirus on all virtual machines (VMs). Customers can install Microsoft Antimalware for Cloud Services and Virtual Machines or another antivirus solution on VMs, and VMs can be routinely reimaged to clean out intrusions that may have gone undetected.

- **Penetration testing**

Microsoft conducts regular penetration testing to improve Azure security controls and processes. Microsoft understands that security assessment is also an important part of our customers' application development and deployment. Therefore, Microsoft has

established a policy for customers to carry out authorized penetration testing on their own—and only their own—applications hosted in Azure.

- **DDoS protection**

Azure has a defense system against Distributed Denial-of-Service (DDoS) attacks on Azure platform services. It uses standard detection and mitigation techniques. Azure's DDoS defense system is designed to withstand attacks generated from outside and inside the platform.

3.1.3 Network protection

Azure networking provides the infrastructure necessary to securely connect VMs to one another and to connect on-premises datacenters with Azure VMs. Because Azure's shared infrastructure hosts hundreds of millions of active VMs, protecting the security and confidentiality of network traffic is critical. In the traditional datacenter model, a company's IT organization controls networked systems, including physical access to networking equipment. In the cloud service model, the responsibilities for network protection and management are shared between the cloud provider and the customer. Customers do not have physical access, but they implement the logical equivalent within their cloud environment through tools such as Guest operating system (OS) firewalls, Virtual Network Gateway configuration, and Virtual Private Networks.

- **Network isolation**

Azure is a multitenant service, meaning that multiple customers' deployments and VMs are stored on the same physical hardware. Azure uses logical isolation to segregate each customer's data from that of others. This provides the scale and economic benefits of multitenant services while rigorously preventing customers from accessing one another's data.

- **Virtual networks**

A customer can assign multiple deployments within a subscription to a virtual network and allow those deployments to communicate with each other through private IP addresses. Each virtual network is isolated from other virtual networks.

- **VPN and ExpressRoute**

Microsoft enables connections from customer sites and remote workers to Azure Virtual Networks using Site-to-Site and Point-to-Site VPNs. For even better performance, customers can use an optional ExpressRoute, a private fiber link into Azure datacenters that keeps their traffic off the Internet.

- **Encrypting communications**

Built-in cryptographic technology enables customers to encrypt communications within

and between deployments, between Azure regions, and from Azure to on-premises datacenters.

3.1.4 Data protection

Azure allows customers to encrypt data and manage keys, and safeguards customer data for applications, platform, system, and storage using three specific methods: encryption, segregation, and destruction.

- **Data isolation**

Azure is a multitenant service, meaning that multiple customers' deployments and virtual machines are stored on the same physical hardware.

- **Protecting data at rest**

Azure offers a wide range of encryption capabilities, giving customers the flexibility to choose the solution that best meets their needs. Azure Disk Encryption is a capability that lets you encrypt your Windows and Linux IaaS virtual machine disks. Azure Disk Encryption leverages the industry-standard BitLocker feature of Windows and the DM-Crypt feature of Linux to provide volume encryption for the OS and the data disks. The solution is integrated with Azure Key Vault to help you control and manage the disk encryption keys and secrets in your key vault subscription, while ensuring that all data in the virtual machine disks are encrypted at rest in your Azure storage.

SQL Database TDE is based on SQL Server's TDE technology, which encrypts the storage of an entire database by using an industry-standard AES-256 symmetric key called the database encryption key. SQL Database protects this database encryption key with a service-managed certificate. All key management for database copying, Geo-Replication, and database restores anywhere in SQL Database is handled by the service.

- **Protecting data in transit**

For data in transit, customers can enable encryption for traffic between their own VMs and end users. Azure protects data in transit, such as between two virtual networks. Azure uses industry-standard transport protocols such as TLS between devices and Microsoft datacenters, and within datacenters themselves.

- **Encryption**

Customers can encrypt data in storage and in transit to align with best practices for protecting confidentiality and data integrity. For data in transit, Azure uses industry-standard transport protocols between devices and Microsoft datacenters and within datacenters themselves. You can enable encryption for traffic between your own virtual machines and end users.

- **Data redundancy**

Customers may opt for in-country storage for compliance or latency considerations or out-of-country storage for security or disaster recovery purposes. Data may be replicated within a selected geographic area for redundancy.

- **Data destruction**

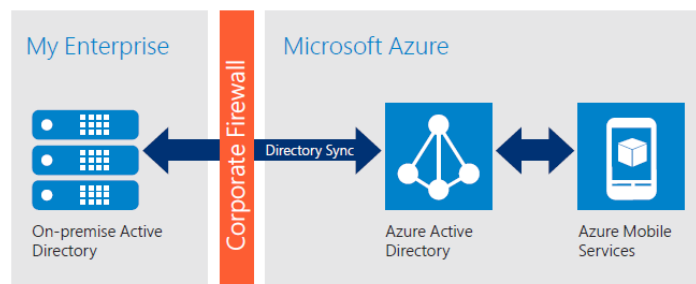
When customers delete data or leave Azure, Microsoft follows strict standards for overwriting storage resources before reuse. As part of our agreements for cloud services such as Azure Storage, Azure VMs, and Azure Active Directory, we contractually commit to specific processes for the deletion of data.

3.1.5 Identity and access

Microsoft has strict controls that restrict access to Azure by Microsoft employees. Azure also enables customers to control access to their environments, data, and applications.

- **Enterprise cloud directory**

Azure Active Directory is a comprehensive identity and access management solution in the cloud. It combines core directory services, advanced identity governance, security, and application access management. Azure Active Directory makes it easy for developers to build policy-based identity management into their applications. Azure Active Directory Premium includes additional features to meet the advanced identity and access needs of enterprise organizations. Azure Active Directory enables a single identity management capability across on-premises, cloud, and mobile solutions.



- **Multi-Factor Authentication**

Microsoft Azure provides Multi-Factor Authentication (MFA). This helps safeguard access to data and applications and enables regulatory compliance while meeting user demand for a simple sign-in process for both on-premises and cloud applications. It delivers strong authentication via a range of easy verification options—phone call, text message, or mobile app notification—allowing users to choose the method they prefer.

- **Access monitoring and logging**

Security reports are used to monitor access patterns and to proactively identify and

mitigate potential threats. Microsoft administrative operations, including system access, are logged to provide an audit trail if unauthorized or accidental changes are made. Customers can turn on additional access monitoring functionality in Azure and use third-party monitoring tools to detect additional threats. Customers can request reports from Microsoft that provide information about user access to their environments.

3.2 Owning and controlling data

Customers will only use cloud providers in which they have great trust. They must trust that the privacy of their information will be protected, and that their data will be used in a way that is consistent with their expectations. Standards and processes that support Privacy by Design principles include the Microsoft Online Services Privacy Statement and the Microsoft Security Development Lifecycle. We then back those protections with strong contractual commitments to safeguard customer data, including offering EU Model Clauses (which provides terms covering the processing of personal information), and complying with international standards. Microsoft uses customer data stored in Azure only to provide the service, including purposes compatible with providing the service. Azure does not use customer data for advertising or similar commercial purposes.

Microsoft was the first major cloud service provider to make contractual privacy commitments that help ensure the privacy protections built into in-scope Azure services are strong. Among the many commitments that Microsoft supports are:

- EU Model Clauses
- US-EU Safe Harbor Framework and the US-Swiss Safe Harbor Program
- ISO/IEC 27018

Access to customer data by Microsoft personnel is restricted. Customer data is only accessed when necessary to support the customer's use of Azure. This may include troubleshooting aimed at preventing, detecting, or repairing problems affecting the operation of Azure and the improvement of features that involve the detection of, and protection against, emerging and evolving threats to the user (such as malware or spam). When granted, access is controlled and logged. Strong authentication, including the use of multifactor authentication, helps limit access to authorized personnel only. Access is revoked as soon as it is no longer needed.

Microsoft believes that customers should control their data whether stored on their premises or in a cloud service. We will not disclose Azure customer data to law enforcement except as a customer directs or where required by law. When governments make a lawful demand for Azure customer data from Microsoft, we strive to be principled, limited in what we disclose, and committed to transparency. In its commitment to transparency, Microsoft regularly publishes a

Law Enforcement Requests Report that discloses the scope and number of requests we receive. Microsoft keeps customers informed about the processes to protect data privacy and security, including practices and policies. Microsoft also provides the summaries of independent audits of services, which helps customers pursue their own compliance.

3.3 Managing compliance and data privacy regulations

Microsoft invests heavily in the development of robust and innovative compliance processes. The Microsoft compliance framework for online services maps controls to multiple regulatory standards. This enables Microsoft to design and build services using a common set of controls, streamlining compliance across a range of regulations today and as they evolve in the future. Microsoft compliance processes also make it easier for customers to achieve compliance across multiple services and meet their changing needs efficiently. Together, security-enhanced technology and effective compliance processes enable Microsoft to maintain and expand a rich set of third-party certifications. These help customers demonstrate compliance readiness to their customers, auditors, and regulators. As part of its commitment to transparency, Microsoft shares third-party verification results with its customers.

Azure meets a broad set of international as well as regional and industry-specific compliance standards, such as ISO 27001, FedRAMP, SOC 1, and SOC 2. Azure's adherence to the strict security controls contained in these standards is verified by rigorous third-party audits that demonstrate Azure services work with and meet world-class industry standards, certifications, attestations, and authorizations.

Azure is designed with a compliance strategy that helps customers address business objectives and industry standards and regulations. The security compliance framework includes test and audit phases, security analytics, risk management best practices, and security benchmark analysis to achieve certificates and attestations. Microsoft Azure offers the following certifications for all in-scope services.

- Content Delivery and Security Association (CDSA)
- Criminal Justice Information Services (CJIS)
- Cloud Security Alliance (CSA) Cloud Controls Matrix
- EU Model Clauses
- US Food and Drug Administration (FDA) Code of Federal Regulations (CFR) Title 21 P 11
- Federal Risk and Authorization Management Program (FedRAMP)
- Family Educational Rights and Privacy Act (FERPA)
- Federal Information Processing Standard (FIPS) Publication 140-2

- Health Insurance Portability and Accountability Act (HIPAA)
- Information Security Registered Assessors Program (IRAP)
- ISO/IEC 27018
- ISO/IEC 27001/27002:2013
- Multi-Level Protection Scheme (MLPS)
- Multi-Tier Cloud Security Standard for Singapore (MTCS SS)
- Payment Card Industry (PCI) Data Security Standards (DSS)
- Service Organization Control (SOC) reporting framework for both SOC 1 Type 2 and SOC 2 Type 2.
- Trusted Cloud Service certification developed by the China Cloud Computing Promotion and Policy Forum (CCCPF)
- UK Government G-Cloud

3.4 Azure Security Center

Security Center helps you prevent, detect, and respond to threats with increased visibility into and control over the security of your Azure resources. It provides integrated security monitoring and policy management across your Azure subscriptions, helps detect threats that might otherwise go unnoticed, and works with a broad ecosystem of security solutions. Security Center delivers easy-to-use and effective threat prevention, detection, and response capabilities that are built in to Azure. Key capabilities are:

Prevent:

- Monitors the security state of your Azure resources
- Defines policies for your Azure subscriptions and resource groups based on your company's security requirements, the types of applications that you use, and the sensitivity of your data
- Uses policy-driven security recommendations to guide service owners through the process of implementing needed controls
- Rapidly deploys security services and appliances from Microsoft and partners

Detect:

- Automatically collects and analyzes security data from your Azure resources, the network, and partner solutions like antimalware programs and firewalls

- Leverages global threat intelligence from Microsoft products and services, the Microsoft Digital Crimes Unit (DCU), the Microsoft Security Response Center (MSRC), and external feeds
- Applies advanced analytics, including machine learning and behavioral analysis

Respond:

- Provides prioritized security incidents/alerts
- Offers insights into the source of the attack and impacted resources
- Suggests ways to stop the current attack and help prevent future attacks

3.5 Microsoft US government cloud

Azure Government is a government-community cloud (GCC) designed to support strategic government scenarios that require speed, scale, security, compliance, and economics for US government organizations. It was developed based on Microsoft's extensive experience delivering software, security, compliance, and controls in other Microsoft cloud offerings such as Azure public, Office 365, Office 365 GCC, Microsoft CRM Online, etc.

In addition, Azure Government is designed to meet the higher level security and compliance needs for sensitive, dedicated, US Public Sector workloads found in regulations such as United States Federal Risk and Authorization Management Program (FedRAMP), Department of Defense Enterprise Cloud Service Broker (ECSB), Criminal Justice Information Services (CJIS) Security Policy, and Health Insurance Portability and Accountability Act (HIPAA).

Azure Government includes the core components of Infrastructure as a Service (IaaS) and Platform as a Service (PaaS). This includes infrastructure, network, storage, data management, identity management, and many other services.

3.6 Microsoft cloud in Germany

Starting in 2016, Microsoft will offer its cloud services Microsoft Azure, Office 365 and Dynamics CRM Online from within German datacenters. That alone wouldn't be really surprising or innovative, but the unique thing about this is that the keys (physical and logical) that control access to customer data in this cloud are held by a German company, Deutsche Telekom's subsidiary T-Systems, which will act as a Data Trustee. So Microsoft will have no access to customer data without approval and supervision by the Data Trustee.

All access rights are handled by a role-based access model, better known as RBAC. Those roles are based on functions (Reader, Owner, etc.) and/or on realms (server, mailboxes, resource

groups, etc.). Microsoft has—in this new model—no rights at all to access customer data. Only for a special purpose like a support call from a customer will a temporary access be granted by the Data Trustee to the Microsoft engineer, and only for the specified area. After that time all access is revoked automatically. Microsoft has no way to grant that access to itself. And of course there is a logging of this process to an area where Microsoft has no access, too. In addition, the Data Trustee is escorting the session and watching the engineer at work.

That RBAC is also in place for physical access to the datacenters. The Data Trustee has to approve the visit and will escort Microsoft or any of its subcontractors at any time during the visit. For all those cases where Microsoft could come in contact with customer data, it needs a reason related to operation of the services (incident, support case), a well-defined area of access, and a well-defined time period, and only then the trustee will grant access.

Customer data is only stored in the German datacenters. Data exchange between the two Azure regions in Germany (Germany Central and Germany Northeast) is handled by a dedicated network line leased from a German provider, just to make sure that no data is accidentally routed outside of Germany. There is no additional replication or backup to other regions; even Azure Active Directory is only replicated between those two German Azure regions.

For encryption and securing data traffic between client applications and cloud servers, Microsoft relies on the nationally recognized certification authority of Bundesdruckerei GmbH, D-TRUST. This ensures customers of Microsoft Cloud Germany that their data is protected by the latest encryption technologies available in the market. With the safety concepts of Bundesdruckerei, users and servers can be reliably authenticated to ensure encrypted traffic.

3.7 Cloud security recommendations for enterprise architects

Although Microsoft is committed to the privacy and security of your data and applications in the cloud, customers must take an active role in the security partnership. Ever-evolving cybersecurity threats increase the requirements for security rigor and principles at all layers for both on-premises and cloud assets. Enterprise organizations are better able to manage and address concerns about security in the cloud when they take a systematic approach. Moving workloads to the cloud shifts many security responsibilities and costs to Microsoft, freeing your security resources to focus on the critically important areas of data, identity, strategy, and governance.

Responsibility	SaaS	PaaS	IaaS	On-prem
Data governance & rights management	Customer	Customer	Customer	Customer
Client endpoints	Customer	Customer	Customer	Customer
Account & access management	Customer	Customer	Customer	Customer
Identity & directory infrastructure	Microsoft	Microsoft	Customer	Customer
Application	Microsoft	Microsoft	Customer	Customer
Network controls	Microsoft	Microsoft	Customer	Customer
Operating system	Microsoft	Microsoft	Customer	Customer
Physical hosts	Microsoft	Microsoft	Microsoft	Customer
Physical network	Microsoft	Microsoft	Microsoft	Customer
Physical datacenter	Microsoft	Microsoft	Microsoft	Customer
	Microsoft	Microsoft	Microsoft	Customer

Your responsibility for security is based on the type of cloud service. The chart summarizes the balance of responsibility for both Microsoft and the customer.

Recommendations for security strategy, governance, and operationalization	
Develop cloud security policies	<p>Policies enable you to align your security controls with your organization's goals, risks, and culture. Policies should provide clear, unequivocal guidance to enable good decisions by all practitioners.</p> <p>Document security policies in enough detail to guide personnel into quick and accurate decisions while adopting and managing cloud services. Ensure you have sufficient detail on policy areas that are well-established and critically important to your security posture.</p> <p>Balance security and usability. Security controls that overly restrict the ability of admins and users to accomplish tasks will be worked around. Build buy-in through both threat education and inclusion in the security design process.</p> <p>Document protocols and processes for performing critically important security tasks such as using administrative credentials, responding to common security events, and recovering from significant security incidents.</p> <p>Embrace shadow IT. Identify the unmanaged use of devices, cloud services, and applications. Identify business requirements that led to their use and the business risk that they bring. Work with business groups to enable required capabilities while mitigating risks</p>
Manage continuous threats	<p>The evolution of security threats and changes requires comprehensive operational capabilities and ongoing adjustments. Proactively manage this risk.</p> <p>Establish operational capabilities to monitor alerts, investigate incidents, initiate remediation actions, and integrate lessons learned.</p> <p>Build external context of threats using available resources such as threat intelligence feeds, Information Sharing and Analysis Centers (ISACs), and other means.</p> <p>Validate your security posture by authorized red team and/or penetration testing activity.</p> <p>White paper: Microsoft Enterprise Cloud Red Teaming White paper: Determined Adversaries and Targeted Attacks</p>

Manage continuous innovation	<p>The rate of capability releases and updates from cloud services requires proactive management of potential security impacts.</p> <p>Define a monthly cadence to review and integrate updates of cloud capabilities, regulatory and compliance requirements, evolving threats, and organizational objectives.</p> <p>Prevent configuration drift with periodic reviews to ensure technologies, configurations, and operational practices stay in compliance with your policies and protocols.</p>
Contain risk by assuming breach	<p>When planning security controls and security response processes, assume an attacker has compromised other internal resources such as user accounts, workstations, and applications. Assume an attacker will use these resources as an attack platform. Modernize your containment strategy by:</p> <p>Identifying your most critical assets such as mission-critical data, applications, and dependencies. Security for these must be at a higher level without compromising usability.</p> <p>Enhancing isolation between security zones by increasing rigor of exception management. Apply threat modeling techniques to all authorized exceptions and analysis of these application data flows, including identities used, data transmitted, application and platform trustworthiness, and ability to inspect interaction.</p> <p>Focus containment within a security zone on preserving integrity of the administrative model rather than on network isolation.</p>

Recommendations for administrative control	
Least privilege admin model	<p>Apply least-privilege approaches to your administrative model, including:</p> <ul style="list-style-type: none"> • Limit the number of administrators or members of privileged groups. • Delegate fewer privileges to accounts. • Provide privileges on demand. • Have existing administrators perform tasks instead of adding additional administrators. • Provide processes for emergency access and rare use scenarios.
Harden security dependencies	<p>Security dependencies include anything that has administrative control of an asset. Ensure that you harden all dependencies at or above the security level of the assets they control. Security dependencies for cloud services</p>

	<p>commonly include identity systems, on-premises management tools, administrative groups and accounts, and workstations where these accounts logon.</p> <p>Microsoft Advanced Threat Analytics</p>
Use strong authentication	<p>Use credentials secured by hardware or Multi-Factor Authentication (MFA) for all identities with administrative privileges. This mitigates risk of stolen credentials being used to abuse privileged accounts.</p> <p>Azure Multi-Factor Authentication</p>
Use dedicated admin accounts and workstations	<p>Separate high-impact assets from highly prevalent Internet browsing and email risks:</p> <ul style="list-style-type: none"> • Use dedicated accounts for privileged administrative roles for cloud services and on-premises dependencies. • Use dedicated, hardened workstations for administration of high-business impact IT assets. • Do not use high privilege accounts on devices where email and web browsing take place. <p>Securing Privileged Access</p>
Enforce stringent security standards	<p>Administrators control significant numbers of organizational assets. Rigorously measure and enforce stringent security standards on administrative accounts and systems. This includes cloud services and on-premises dependencies such as Active Directory, identity systems, management tools, security tools, administrative workstations, and associated operating systems.</p>
Monitor admin accounts	<p>Closely monitor the use and activities of administrative accounts. Configure alerts for activities that are high impact as well as for unusual or rare activities.</p> <p>White paper: Microsoft Azure Security and Audit Log Management</p>
Educate and empower admins	<p>Educate administrative personnel on likely threats and their critical role in protecting their credentials and key business data. Administrators are the gatekeepers of access to many of your critical assets. Empowering them with this knowledge will enable them to be better stewards of your assets and security posture.</p>

Recommendations for data protection

Establish information protection priorities	<p>The first step to protecting information is identifying what to protect. Develop clear, simple, and well-communicated guidelines to identify, protect, and monitor the most important data assets anywhere they reside.</p> <p>Trustworthy Computing: Data governance</p>
Protect High Value Assets (HVAs)	<p>Establish the strongest protection for assets that have a disproportionate impact on the organization's mission or profitability. Perform stringent analysis of HVA lifecycle and security dependencies, and establish appropriate security controls and conditions.</p>
Find and protect sensitive assets	<p>Identify and classify sensitive assets. Define the technologies and processes to automatically apply security controls.</p> <p>Azure Rights Management</p> <p>Azure Key Vault</p> <p>Always Encrypted (Database Engine)</p>
Set organizational minimum standards	<p>Establish minimum standards for trusted devices and accounts that access any data assets belonging to the organization. This can include device configuration compliance, device wipe, enterprise data protection capabilities, user authentication strength, and user identity.</p>
Establish user policy and education	<p>Users play a critical role in information security and should be educated on your policies and norms for the security aspects of data creation, classification, compliance, sharing, protection, and monitoring.</p>

Recommendations for user identity and device security

Use strong authentication	<p>Use credentials secured by hardware or Multi-Factor Authentication (MFA) for all identities to mitigate the risk that stolen credentials can be used to abuse accounts.</p> <ul style="list-style-type: none"> User identities hosted in Azure Active Directory (Azure AD). On-premises accounts whose authentication is federated from on-premises Active Directory. <p>Azure Multi-Factor Authentication</p>
Manage trusted and compliant devices	<p>Establish, measure, and enforce modern security standards on devices that are used to access corporate data and assets. Apply configuration standards</p>

	and rapidly install security updates to lower the risk of compromised devices being used to access or tamper with data.
Educate, empower, and enlist users	<p>Users control their own accounts and are on the front line of protecting many of your critical assets. Empower your users to be good stewards of organizational and personal data. At the same time, acknowledge that user activities and errors carry security risks that can be mitigated but never completely eliminated. Focus on measuring and reducing risks from users.</p> <ul style="list-style-type: none"> • Educate users on likely threats and their role in protecting business data. • Increase adversary cost to compromised user accounts. • Explore gamification and other means of increasing user engagement.
Monitor for account and credential abuse	<p>One of the most reliable ways to detect abuse of privileges, accounts, or data is to detect anomalous activity of an account.</p> <ul style="list-style-type: none"> • Identify activity that is normal and physically possible. Alert on unusual activity to enable rapid investigation and response. • For accounts in Azure AD, use the integrated analytics to detect unusual activity. <p>White paper: Microsoft Azure Security and Audit Log Management</p>

Recommendations for application security

Secure applications that you acquire	<ul style="list-style-type: none"> • Review the security development processes and operational practices of vendors before acquiring applications. Build this into your acquisition process. • Follow security configuration guidance and recommendations provided by the vendor for the application. • Apply all vendor security updates as rapidly as your testing requirements allow. Be sure to update middleware and dependencies installed with the applications. • Discontinue your use of software before it reaches end of support status.
Follow the Security Development Lifecycle (SDL)	<p>Software applications with source code you develop or control are a potential attack surface. These include PaaS apps, PaaS apps built from sample code in Azure (such as WordPress sites), and apps that interface with Office 365. Follow code security best practices in the Microsoft Security</p>

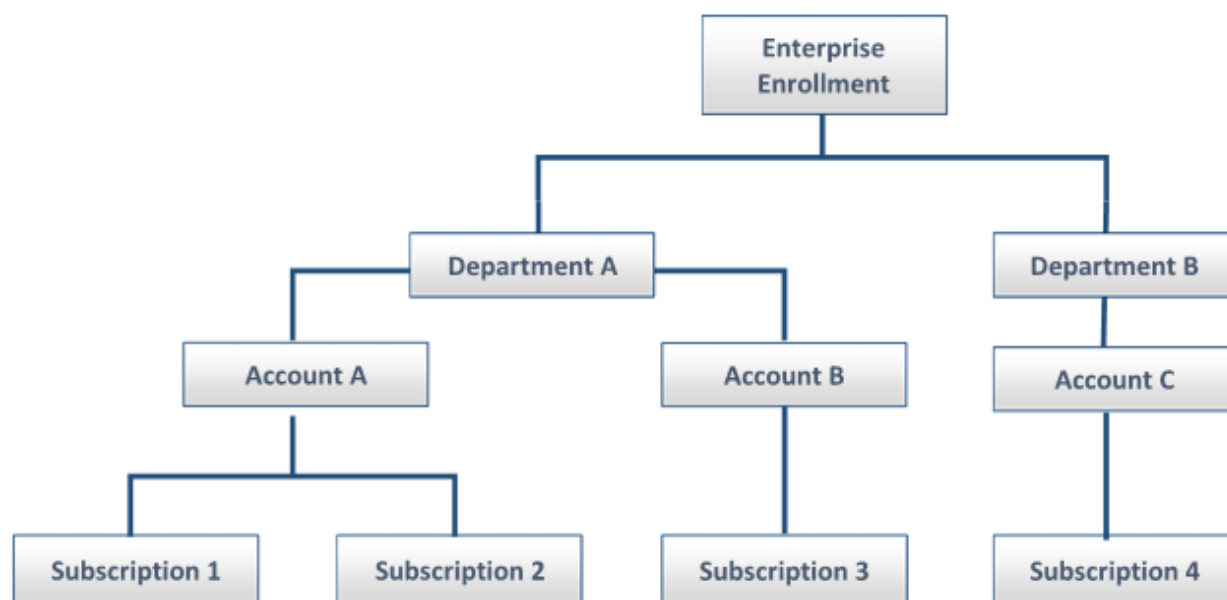
	<p>Development Lifecycle (SDL) to minimize vulnerabilities and their security impact.</p> <p>See www.microsoft.com/sdl.</p>
--	---

Recommendations for network connectivity	
Update your network security strategy and architecture for cloud computing	<p>Ensure your network architecture is ready for the cloud by updating your current approach or taking the opportunity to start fresh with a modern strategy for cloud services and platforms. Align your network strategy with your:</p> <ul style="list-style-type: none"> • Overall security strategy and governance • Containment model and identity strategy • Cloud services capabilities and constraints <p>Your design should address securing communications:</p> <ul style="list-style-type: none"> • Inbound from the Internet • Between VMs in a subscription • Across subscriptions • To and from on-premises networks • From remote administration hosts
Optimize with cloud capabilities	<p>Cloud computing offers uniquely flexible network capabilities as topologies are defined in software. Evaluate the use of these modern cloud capabilities to enhance your network security auditability, discoverability, and operational flexibility.</p>
Manage and monitor network security	<p>Ensure your processes and technology capabilities are able to distinguish anomalies and variances in configurations and network traffic flow patterns. Cloud computing utilizes public networks, allowing rapid exploitation of misconfigurations that should be avoided or rapidly detected and corrected.</p> <ul style="list-style-type: none"> • Closely monitor and alert on exceptions. • Apply automated means to ensure your network configuration remains correct and unusual traffic patterns are detected.

Recommendations for operating system and middleware	
Virtual operating system	<p>Secure the virtual host operating system (OS) and middleware running on virtual machines. Ensure that all aspects of the OS and middleware security meet or exceed the level required for the host, including:</p> <ul style="list-style-type: none"> • Administrative privileges and practices • Software updates for OS and middleware • Security Configuration Baseline • Use of Group Policy Objects (GPOs) • Installation methods and media • Use of scheduled tasks • Antimalware and intrusion detection/prevention • Host firewall and IPsec configurations • Event log configuration and monitoring
Virtual OS management tools	<p>System management tools have full technical control of the host operating systems (including the applications, data, and identities), making these a security dependency of the cloud service. Secure these tools at or above the level of the systems they manage. These tools typically include:</p> <ul style="list-style-type: none"> • Configuration Management • Operations Management and Monitoring • Backup • Security Update and Patch Management <p>Microsoft Cloud Services and Network Security</p> <p>Microsoft Azure Security blog</p> <p>Operations Management Suite</p>

4 Azure enterprise administration

The Azure Enterprise Agreement portal allows large enterprise customers of Azure to manage Azure subscriptions and associated licensing information from a central portal. Enterprise Agreement (EA) customers can add Azure to their EA by making an upfront monetary commitment to Azure. That commitment is consumed throughout the year by using any combination of the wide variety of cloud services Azure offers from its global datacenters. Within a given enterprise enrollment, Microsoft Azure has several roles that individuals play. The Enterprise Administrator has the ability to add or associate Accounts and Departments to the Enrollment, can view usage data across all Accounts and Departments, and is able to see the monetary commitment balance associated to the Enrollment. There is no limit to the number of Enterprise Administrators on an Enrollment.



Departments can be leveraged if an additional level to structure the Accounts and Subscriptions is needed. Cost center and Start/End date can be added as an attribute to the Department. Department Administrators can manage Department properties, manage accounts under the department they administer, download usage details, and view monthly Usage and Charges associated to their Department if the Enterprise Administrator has granted permission to do so. The Account Owner can add Subscriptions for their Account, update the Service Administrator and Co-Administrator for an individual Subscription, view usage data for their Account, and view Account charges if the Enterprise Administrator has provided access. Account Owners will not have visibility of the monetary commitment balance unless they also have Enterprise Administrator rights. The Service Administrator and up to 200 Co-Administrators per Subscription have the ability to access and manage Subscriptions and development projects


within the classic Azure Management Portal. Service Administrators do not have access to the Enterprise Portal unless they also have one of the other two roles. The Resource Group Administrators manage a group of resources within a subscription that collectively provide a service and share a lifecycle: single project or service focused.

The primary tools that are used by these roles are:

- Enterprise Administrator → <https://ea.azure.com>
- Departmental Administrator → <https://account.windowsazure.com>
- Account Owner → <https://account.windowsazure.com>
- Service Administrator → <https://manage.windowsazure.com>
- Co-Administrator → <https://manage.windowsazure.com>
- Resource Group Administrator → <https://portal.azure.com>

4.1 Understanding Azure subscriptions

Initially, a subscription was the administrative security boundary of Microsoft Azure. With the advent of the Azure Resource Management (ARM) model, a subscription now has two administrative models: Azure Service Management and Azure Resource Management. With ARM, the subscription is no longer needed as an administrative boundary. ARM provides a more granular Role-Based Access Control (RBAC) model for assigning administrative privileges at the resource level. Using RBAC, you can segregate duties within your team and grant only the amount of access to users that they need to perform their jobs. Access is granted by assigning the appropriate RBAC role to users, groups, and applications at a certain scope. The scope of a role assignment can be a subscription, a resource group, or a single resource. A role assigned at a parent scope also grants access to the children contained within it. For example, a user with access to a resource group can manage all the resources it contains, like websites, virtual machines, and subnets.



The RBAC role that is assigned dictates what resources the user, group, or application can manage within that scope. Azure RBAC has three basic roles that apply to all resource types:

- Owner has full access to all resources including the right to delegate access to others.
- Contributor can create and manage all types of Azure resources but can't grant access to others.
- Reader can view existing Azure resources.

The rest of the RBAC roles in Azure allow management of specific Azure resources. For example, the Virtual Machine Contributor role allows users to create and manage virtual machines. It does not give them access to the virtual network or the subnet that the virtual machine connects to.

A subscription additionally forms the billing unit. Services charges are accrued to the subscription. As part of the new Azure Resource Management model, it is also possible to roll up costs. A standard naming convention for Azure resource object types can be used to manage billing across projects teams, business units, or other desired view. See section 9.4 for further details.

A subscription is also a logical limit of scale by which resources can be allocated. These limits include hard and soft caps of various resource types (see <https://azure.microsoft.com/en-us/documentation/articles/azure-subscription-service-limits/>). Scalability is a key element for understanding how the subscription strategy will account for growth as consumption increases. If you want to raise the limit above the Default Limit, you can open an [online customer support request](#) at no charge.

Every Azure subscription has a trust relationship with an Azure Active Directory instance (see section 6 for further details). This means that it trusts that directory to authenticate users, services, and devices. Multiple subscriptions can trust the same directory, but a subscription trusts only one directory.

This trust relationship that a subscription has with a directory is unlike the relationship that a subscription has with all other resources in Azure (websites, databases, and so on), which are more like child resources of a subscription. If a subscription expires, then access to those other resources associated with the subscription also stops. But the directory remains in Azure, and you can associate another subscription with that directory and continue to manage the directory users.

As a best practice, you should sign up for Azure as an organization and use a work or school account to manage resources in Azure. Work or school accounts are preferred because they can be centrally managed by the organization that issued them, they have more features than Microsoft accounts, and they are directly authenticated by Azure Active Directory. The same account provides access to other Microsoft online services that are offered to businesses and organizations, such as Office 365 or Microsoft Intune. If you already have an account that you use with those other properties, you likely want to use that same account with Azure.

The important point here is that Azure subscription admins and Azure AD directory admins are two separate concepts. Azure subscription admins can manage resources in Azure. Directory admins can manage properties in the directory. A person can be in both roles, but this isn't required.

4.2 Managing Azure subscriptions

In an Enterprise Environment it is key to set up Azure subscriptions in a way that ensures they support the requirements and fulfil the needs for reporting, segregation, and management today and the future. It is also important to minimize migration of resources between subscriptions because of subscription reorganizations. There are several motivations for using multiple Azure subscriptions. The most common ones are:

- **Project-based billing and chargeback**

There is a desire that individual projects get their own Azure bills. Today within Azure the lowest level of cost aggregation is at the subscription level. Overcoming this constraint is possible by using third-party tools with additional billing and chargeback capabilities or by creating an individual solution with the APIs described in section 9.

- **Reuse of shared infrastructure**

Some applications and services will be dependent on components of shared infrastructure. The most popular scenario is sharing a common VPN to on-premises infrastructure. Today Azure imposes a constraint that a VPN is tied to a specific Virtual Network, which in turn is allocated to a specific Subscription. It is possible to connect different Azure Virtual Networks together with additional Site to Site VPNs. The Azure Virtual Networks can be in the same or in different subscriptions.

One ExpressRoute circuit can connect multiple Azure Virtual Networks across multiple subscriptions as long as the location of the Virtual Networks is connected with the ExpressRoute circuit. These constraints drive project teams in a direction for a shared subscription model in a lot of cases.

- **Security least privilege**

A subscription is the security boundary such that an administrator on a subscription can modify any resources within that subscription. If subscriptions are shared across teams, then these administrators have greater rights than they need to perform their role, increasing the security risk profile. Role-Based Access Control (RBAC) gives the possibility to assign roles and rights on Resource Group and Resource level. There are currently more than 20 Built-in roles available that can be assigned to users and groups to allow a granular assignment of permissions within an Azure subscription. This reduces the amount of required subscriptions significantly. For creating additional custom RBAC roles please refer to <https://azure.microsoft.com/en-us/documentation/articles/role-based-access-control-custom-roles>.

One of the most critical items in the process of designing a subscription is assessing your current environment and needs. Specifically, it is important to have a thorough understanding of the following aspects:

Identify business requirements

- Availability
- Recoverability
- Performance

Identify technical requirements

- Is network connectivity a shared resource or dedicated to single use or group?
- Are there Active Directory requirements?
- Do you need to consider clustering, identity, or management tools?

Security requirements

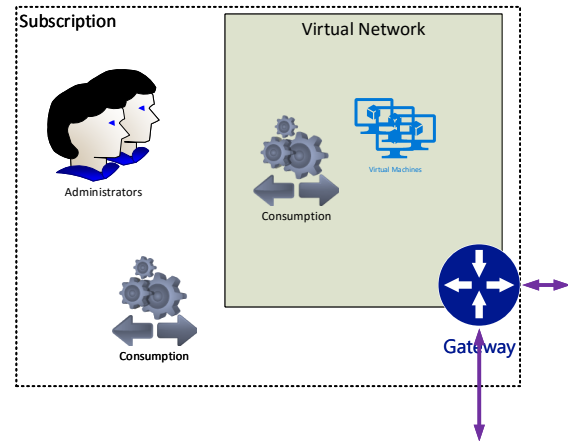
- Who are the subscription administrators?
- Are the appropriate network connectivity and identity requirements being deployed?
- Have you implemented a least privilege administrative model?

Scalability requirements

- What are the growth plans?
- How will limited resources be allocated?
- How will the model evolve over time considering additional users, shared access, and resource limits?

Adding network connectivity (whether using a site-to-site VPN or a dedicated ExpressRoute connection) brings additional considerations to the subscription requirements discussion. For more information about network design, see section 5 in this document.

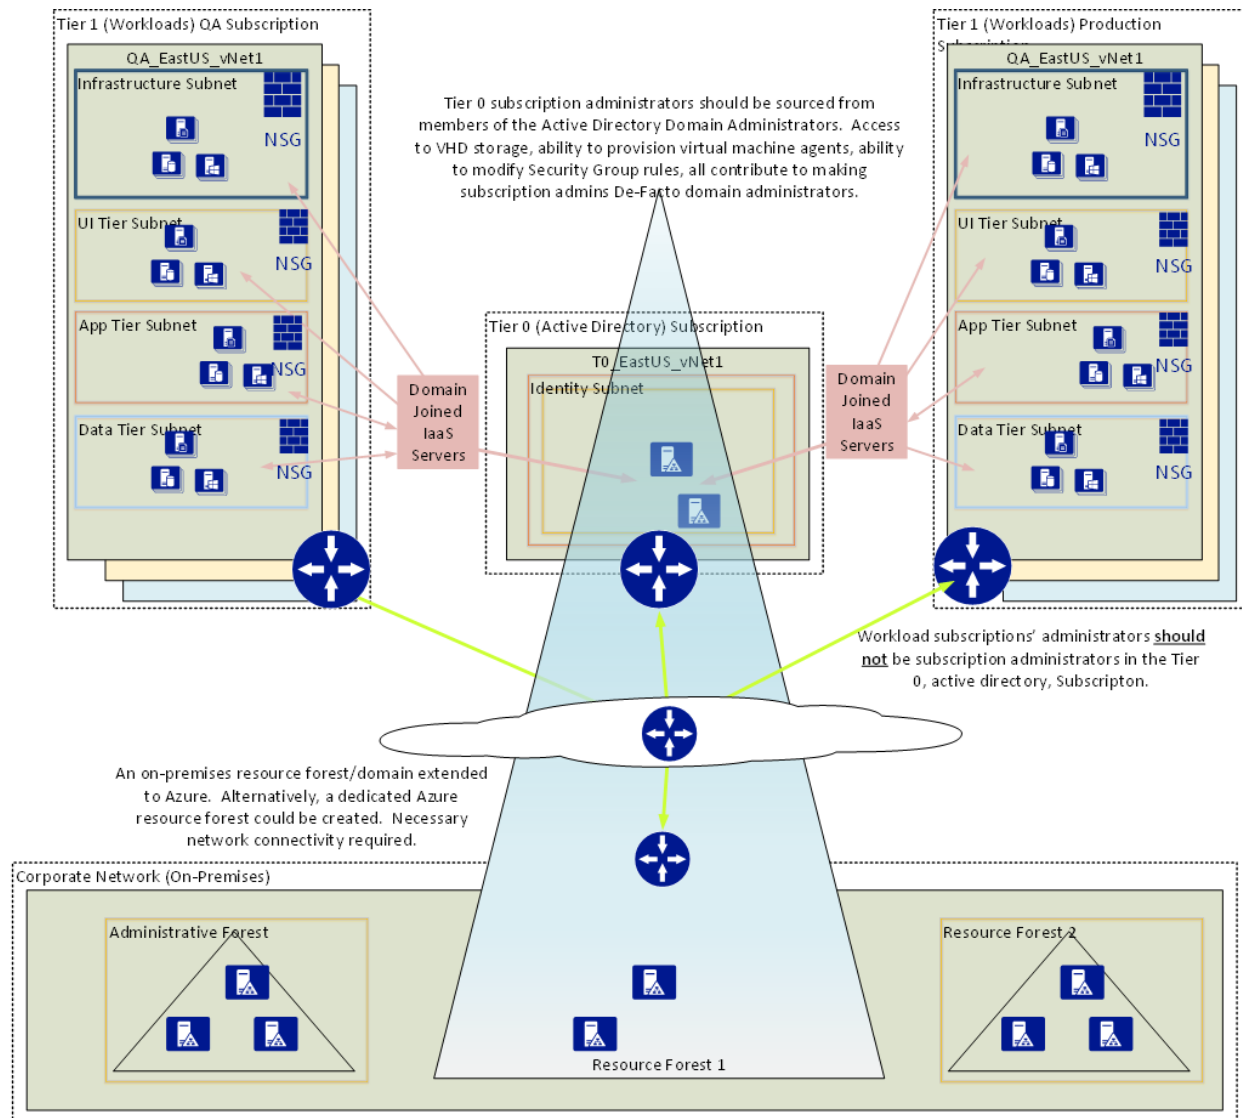
The subscription is a required container to hold a virtual network, and often networking is a shared resource within an enterprise. Site-to-site VPNs and ExpressRoute circuits require defining IP address ranges that do not overlap with on-premises ranges. Site-to-site VPN connectivity requires setting up and configuring a public-facing gateway and VPN services at the corporate edge. ExpressRoute connectivity is through a private connection from an on-premises datacenter to Azure through a service provider's private network. Routing and firewall configurations are typically necessary when enabling connectivity.



If multiple virtual networks are to share a single enterprise ExpressRoute connection, essentially there is no network isolation between those networks. In this case, any separation the subscription design may try to define is eliminated and must be achieved through subnet layer Network Security Groups (NSGs). When the virtual networks are attached to the same ExpressRoute circuit, they are essentially a single routing domain. A subscription hosting only PaaS services could have no virtual network at all, and the design limitations discussed above would not apply.

The following diagram shows a robust enterprise Azure enrollment. There are multiple subscriptions, one of which is a "Tier 0" subscription used to host shared resources such as domain controllers and other sensitive roles when extending an on-premises Active Directory forest to Azure.

This is configured as a separate subscription to ensure that only administrators with domain administrator level privileges are able to exert administrative control over these sensitive servers through Azure subscriptions, while still allowing server administrators to manage virtual machines in other subscriptions.



QA and production networks share the same dedicated ExpressRoute circuit to on-premises resources. They are separated into distinct subscriptions to allow separation of access and to allow the QA subscription to scale on its own without impacting production.

This model will scale based on need. Second, third, and subsequent QA and production subscriptions can be added to this design without significant impact on operations. Those subscriptions can be managed by the project teams they belong to. The same scalability applies to network bandwidth—the circuit can be used until its limits are reached without any artificial limitations forcing additional purchases.

A typical subscription model will be based on a mixed model of Shared subscriptions and Project subscriptions (or business department subscriptions) driven by particular project requirements.

4.3 Defining naming conventions

When naming the Microsoft Azure subscription, it is a recommend practice to be verbose. Try using the following format or a format that has been agreed on by the stakeholders of the company.

<Company> <Department (optional)> <Product Line (optional)> <Environment>

- **Company**, in most cases, would be the same for each subscription. However, some companies may have child companies within the organizational structure. These companies may be managed by a central IT group, in which case they could be differentiated by having both the parent company name and child company name.
- **Department** is a name within the organization where a group of individuals work. This item within the namespace is optional. This is because some companies may not need to drill into such detail due to their size. The company may want to use a different identifier.
- **Product line** is a specific name for a product or function that is performed from within the department. As with the department namespace, this area is optional and can be swapped out as needed.
- **Environment** is the name that describes the deployment lifecycle of the applications or services, such as Dev, Lab, or Prod.

What you are trying to accomplish with a naming convention is to put together a meaningful name about the particular subscription and how it is represented within the company. Many organizations will have more than one subscription, which is why it is important to have a naming convention and use it consistently when creating subscriptions.

4.4 Recommendations for Azure enterprise administration

Recommendations for Azure enterprise administration	
Limit the number of administrative users	Assign a minimum number of users as Subscription Administrators and/or Co-administrators.
Use Role-Based Access	Use Azure Resource Management RBAC whenever possible to control the amount of access that administrators have, and log what changes are made to the environment.
Use work accounts	You should sign up for Azure as an organization and use a work or school account to manage resources in Azure. Do not allow the use of existing personal Microsoft Accounts.

Define naming conventions	Assign meaningful names to your Azure subscriptions according to defined naming conventions.
Use Tier 0 subscription	Use Tier 0 subscription to host shared resources, such as domain controllers and other sensitive roles, and limit the privileges to access it.
Use project subscriptions	Use decentralized project subscriptions. Delegate management of those subscriptions to the responsible project teams.
Separate production from QA	Separate QA environments into distinct subscriptions to allow separation of access and to allow the QA subscription to scale on its own without impacting production.

5 Integrating Azure into the corporate network

Within Azure, there is the concept of virtual networks, subnets within the virtual networks, and the network gateways that allow connectivity between virtual networks and on-premises networks.

Virtual networks can be used to allow isolated network communication within the Azure environment or establish cross-premises network communication between an organization's network infrastructure and Azure. By default, when virtual machines are created and connected to Azure Virtual Network, they are allowed to route to any subnet within the virtual network, and outbound access to the Internet is provided by Azure's Internet connection.

A fundamental first step in creating services within Microsoft Azure is establishing a Virtual Network. To establish a virtual private network within Azure, you must create a minimum of one virtual network. Each virtual network must contain an IP address space and a minimum of one subnet that leverages all or part of the virtual network address space.

5.1 Choosing the right connectivity option

To establish remote network communications to on-premises or other virtual networks, a gateway subnet must be allocated for the virtual network and a virtual network gateway must be added to it. To enable cross-premises connectivity, a Virtual Network must attach a virtual network gateway.

Currently, there are three types of gateways that can be deployed:

- Static routing gateway for Site-to-Site (S2S) VPN connections (basic, standard, and high performance)
- Dynamic routing gateway for Site-to-Site (S2S) VPN connections (basic, standard, and high performance)
- Dynamic Routing ExpressRoute gateway (standard and high performance)

The type of gateway determines the cross-premises connectivity capabilities, the performance, and the features that are offered. Static and dynamic gateways are used when establishing Point-to-Site (P2S) and Site-to-Site (S2S) VPN connections where the cross-premises connectivity leverages the Internet for the transport path. ExpressRoute gateways are designed for high-speed, private, cross-premises connectivity where the traffic flows across dedicated circuits and not the Internet.

A static routing gateway uses policy-based VPNs. Policy-based VPNs encrypt and route packets through an interface based on a customer-defined policy. Static gateways are for establishing low-cost connections to a single virtual network in Azure.


Dynamic routing gateways use route-based VPNs. Route-based VPNs depend on a tunnel interface specifically created for forwarding packets. Any packet arriving at the tunnel interface is forwarded through the VPN connection. Dynamic gateways are used to establish low-cost connections to an on-premises environment or to connect multiple virtual networks for routing purposes in Azure. In addition, it supports Border Gateway Protocol (BGP) Routing (see <https://azure.microsoft.com/en-us/documentation/articles/role-based-access-control-custom-roles/>).

ExpressRoute gateways are always dynamic routing gateways that support BGP routing protocols. ExpressRoute gateways are used for connecting on-premises environments to Azure over high-speed private connections.


For Site-to-Site gateways, an IPsec/IKE VPN tunnel is created between the virtual networks and the on-premises sites by using Internet Key Exchange (IKE) protocol handshakes. For ExpressRoute, the gateways advertise the prefixes by using the BGP in your virtual networks via the peering circuits. The gateways also forward packets from your ExpressRoute circuits to your virtual machines inside your virtual networks.

Each gateway has a limited number of other gateway connections that it can establish. The connection model between gateways dictates how far you can route within Azure. There are three distinct models that you can leverage to connect multiple virtual networks to one another:


Mesh



Hub and Spoke



Daisy-Chain



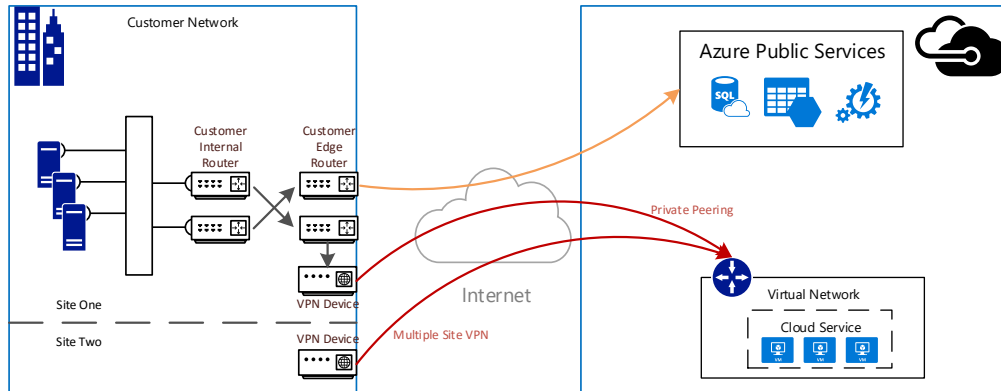
In the Mesh approach, every virtual network can talk to every other virtual network with a single hop. Therefore, this approach does not require you to define multiple hop routing. Challenges with this approach include the rapid consumption of gateway connections, which limits the size of the virtual network routing capability.

In the Hub and Spoke approach a virtual machine on vNet1 will be able to communicate to a virtual machine on vNet2, vNet3, vNet4, or vNet5. A virtual machine on vNet2 could talk to virtual machines on vNet1, but not a virtual machine on vNet3, vNet4, or vNet5. This is due to the default single hop isolation of the virtual network in this configuration.

In a Daisy-Chain approach, a virtual machine on vNet1 can communicate to a virtual machine on vNet2, but not vNet3, vNet4, or vNet5. A virtual machine on vNet2 could talk to virtual machines on vNet1 and vNet3. The same virtual network single hop isolation applies.

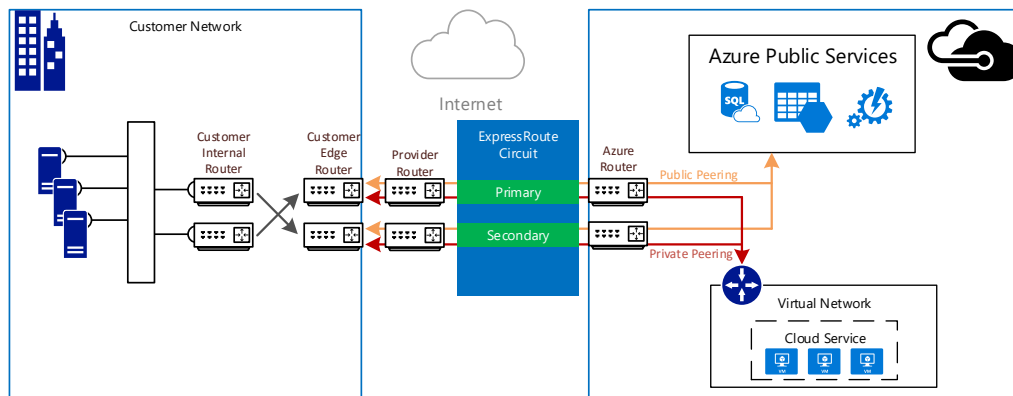
Azure supports two types of connectivity options to connect customers' networks to Azure virtual networks: Site-to-Site VPN and ExpressRoute. Although Point-to-Site is another viable connectivity option, it is client-focused and is not specific to this section.

Site-to-Site VPN connections use VPN devices over public Internet connections to create a path to route traffic to a virtual network in a customer subscription. Traffic to the virtual network flows across an encrypted VPN connection, while traffic to the Azure public services flows over the Internet. It is not possible to create a Site-to-Site VPN connection that provides direct connectivity to the public Azure services via a public peering path. To provide multiple VPN connections to the virtual network, you must use multiple VPN devices connected to different sites. These relationships are depicted in the following diagram:



ExpressRoute connections use routers and private network paths to route traffic to Azure Virtual Network and, optionally, to the Azure public services. Private connections are made through a network provider by establishing an ExpressRoute circuit with a selected provider. The customer's router is connected to the provider's router, and the provider creates the ExpressRoute circuit to connect to the Azure Routers.

When the circuit is created, VLANs can be created that allow separate paths to the private peering network to link to virtual networks and to the public peering network to access Azure public services.



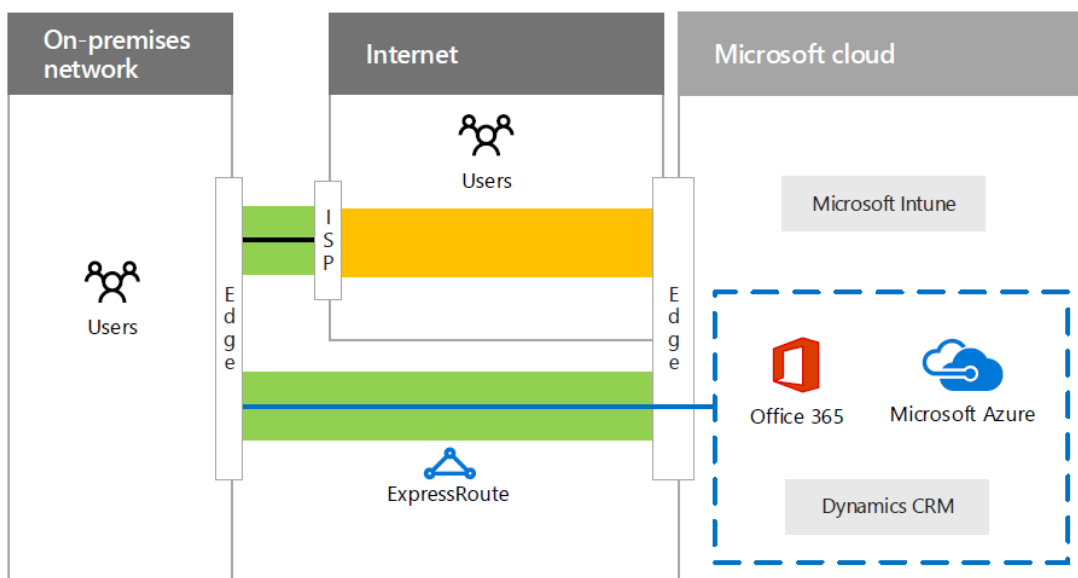
The best type of connection is depending on the detailed requirements. Nevertheless, we can say that enterprise customers tend to use ExpressRoute to fulfill their security and bandwidth requirements.

5.1.1 Using ExpressRoute

With an Internet connection, the only part of the traffic path to the Microsoft cloud that you can control (and have a relationship with the service provider) is the link between your on-premises network edge and your Internet service provider (ISP). The path between your ISP and the

Microsoft cloud edge is a best-effort delivery system subject to outages, traffic congestion, and monitoring by malicious users (shown in yellow).

With an ExpressRoute connection, you now have control, through a relationship with your service provider, over the entire traffic path from your edge to the Microsoft cloud edge. This connection can offer predictable performance and a 99.9 percent uptime SLA. With a dedicated path to the edge of the Microsoft cloud, your performance is not subject to Internet provider outages and spikes in Internet traffic. You can determine and hold your providers accountable to a throughput and latency SLA to the Microsoft cloud. Traffic sent over your dedicated ExpressRoute connection is not subject to Internet monitoring or packet capture and analysis by malicious users. It is as secure as using Multiprotocol Label Switching (MPLS)-based WAN links. With wide support for ExpressRoute connections by exchange providers and network service providers, you can obtain up to a 10 Gbps link to the Microsoft cloud.



Azure ExpressRoute lets you extend your on-premises networks into the Microsoft cloud over a dedicated private connection facilitated by a connectivity provider. With ExpressRoute, you can establish connections to Microsoft cloud services, such as Microsoft Azure, Office 365, and CRM Online. Connectivity can be from an any-to-any (IP VPN) network, a point-to-point Ethernet network, or a virtual cross-connection through a connectivity provider at a co-location facility.

You can create a connection between your on-premises network and the Microsoft cloud in three different ways:

- **Co-located at a cloud exchange**

If you are co-located in a facility with a cloud exchange, you can order virtual cross-connections to the Microsoft cloud through the co-location provider's Ethernet

exchange. Co-location providers can offer either Layer 2 cross-connections, or managed Layer 3 cross-connections between your infrastructure in the co-location facility and the Microsoft cloud.

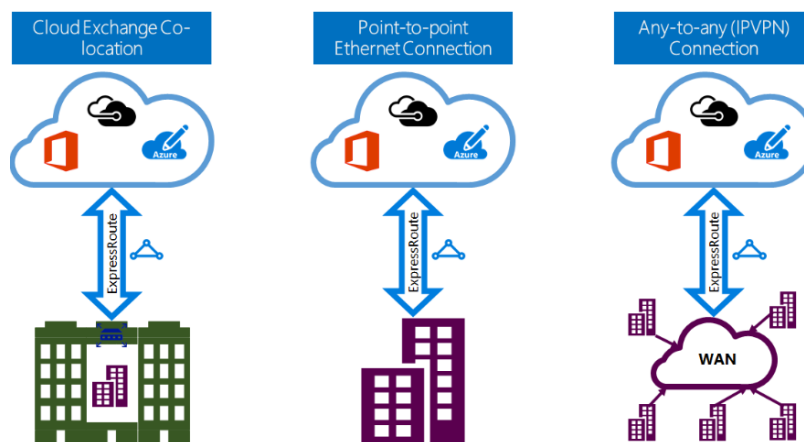
- **Point-to-point Ethernet connections**

You can connect your on-premises datacenters/offices to the Microsoft cloud through point-to-point Ethernet links. Point-to-point Ethernet providers can offer Layer 2 connections, or managed Layer 3 connections between your site and the Microsoft cloud.

- **Any-to-any (IPVPN) networks**

You can integrate your WAN with the Microsoft cloud. IPVPN providers (typically MPLS VPN) offer any-to-any connectivity between your branch offices and datacenters. The Microsoft cloud can be interconnected to your WAN to make it look just like any other branch office. WAN providers typically offer managed Layer 3 connectivity. ExpressRoute capabilities and features are all identical across all of the above connectivity models.

Connectivity providers can offer one or more connectivity models. You can work with your connectivity provider to pick the model that works best for you.

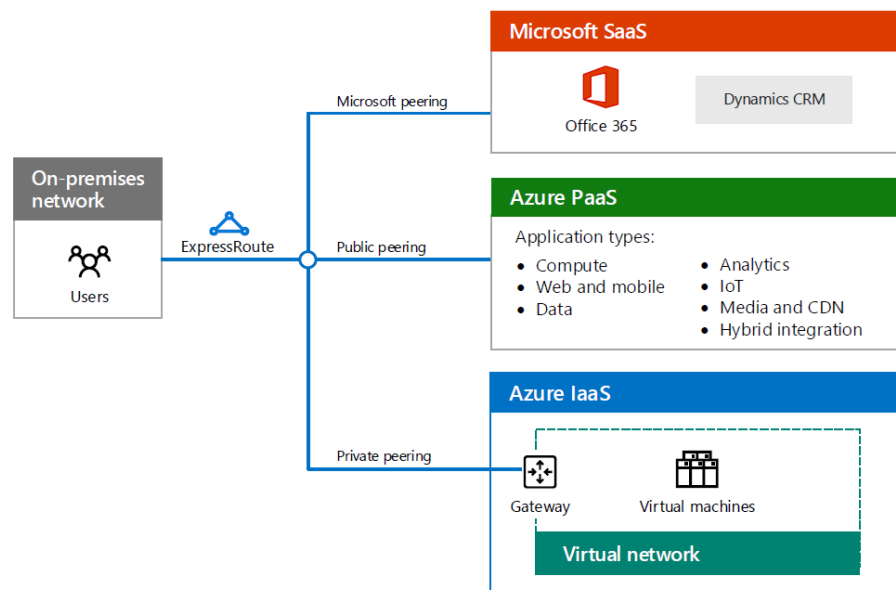


The connections between the customer's network edge and the provider's network edge are redundant as are the connections from the provider's edge to the Azure edge.

From the provider to the Azure edge, you can have private peering connections to customer virtual networks and public peering connections to the Azure PaaS services, such as Azure SQL Database. Pricing models are MeteredData and UnlimitedData. Any provider can provide speeds from 50 Mbps to 10 Gbps.

	MeteredData	UnlimitedData
Bandwidth	50, 100, 200, 500, 1000, 2000, 5000, 10000 Mbps	50, 100, 200, 500, 1000, 2000, 5000, 10000 Mbps
Route management	Varies by provider	Varies by provider
Azure circuit costs	Based on consumption	Unlimited ingress and egress allocation included in monthly fee

Establishing a connection to the public peering network allows virtual machines on Azure Virtual Networks and on-premises systems to leverage the ExpressRoute circuit to connect to Azure PaaS services on the public peering network without traversing the Internet. Establishing a public peering connection is an optional configuration step for an ExpressRoute circuit. When the public peering connection is established, the routes for all the Azure datacenters worldwide are published to the edge router. This directs traffic to the Azure services instead of going out to the Internet.



ExpressRoute connectivity and pricing is made of two components: the service connection costs (Azure) and the authorized carrier costs (telco partner). Customers are charged by Azure for the ExpressRoute monthly access fee, and potentially an egress traffic fee based on the type and performance of the ExpressRoute connection. Customers also have costs associated with the selected provider, which is typically composed of the circuit connection and monthly traffic fees.

From an Azure perspective, an UnlimitedData connection is an inclusive plan where customers are charged a monthly fee and get unlimited ingress and egress traffic. Fees associated with

MeteredData connections include a monthly service charge and traffic egress charges per each GB of traffic transferred based on the zone.

ExpressRoute circuits are created within a subscription. In order to allow an ExpressRoute circuit that was created in one subscription to connect to a virtual network in another subscription, the circuit owner must authorize the connection. This is influencing the subscription management as outlined in section 4.

Enterprise customers with a global business should consider ExpressRoute Premium, which is an add-on package that allows an increase in the number of BGP routes, increases the number of virtual networks per ExpressRoute circuit, and most important allows global connectivity.

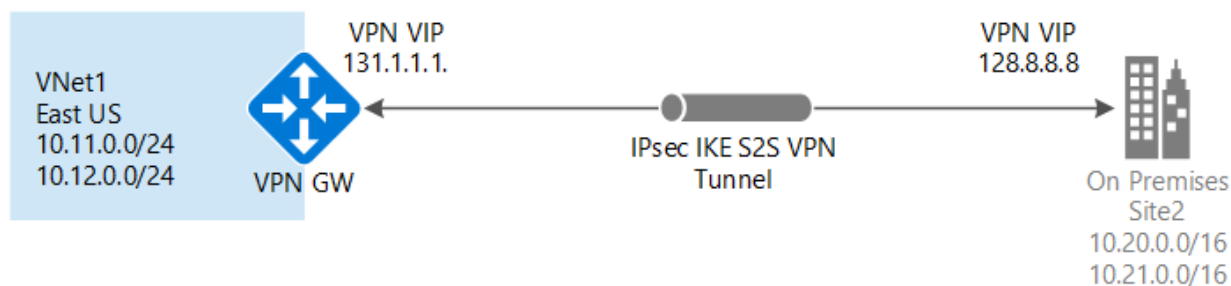
Premium features are:

- Increased route limits for public and private peering (from 4,000 routes to 10,000 routes).
- Global connectivity for services. An ExpressRoute circuit created in any region (excluding China and government clouds) will have access to resources across any other region in the world. For example, a virtual network created in West Europe can be accessed through an ExpressRoute circuit provisioned in the West US region.
- Increased number of virtual network links per ExpressRoute circuit (from 10 to a larger limit, depending on the bandwidth of the circuit).

5.1.2 Using Site-to-Site VPN

Azure Site-to-Site (S2S) connectivity allows low-cost connections from customer locations to Azure private peering networks. S2S leverages the Internet for transport and IPsec encryption to protect the data flowing across the connection. Prerequisites are:

- Public-facing IPv4 address for the on-premises VPN device that is not behind a NAT
- Compatible hardware VPN device or RRAS



5.2 Protecting virtual networks

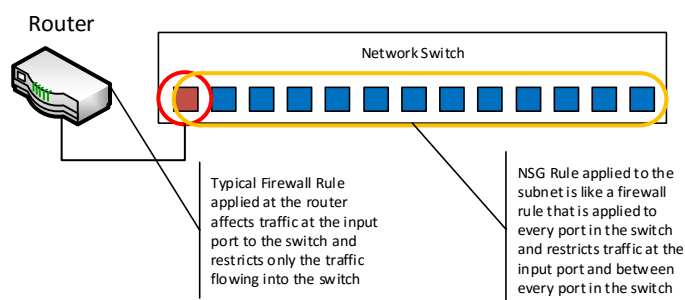
Azure offers multiple different capabilities that could be combined to protect the network.

5.2.1 Network Security Groups

A Network Security Group is a top-level object that is associated with a subscription. It can be used to control traffic to one or more virtual machine instances in the virtual network. A Network Security Group contains access control rules that allow or deny traffic to virtual machine instances. The rules of a Network Security Group can be changed at any time, and changes are applied to all associated instances.

Network Security Groups are similar to firewall rules in that they provide the ability to control the inbound and outbound traffic to a subnet, a virtual machine, or virtual network adapter. Network Security Groups allow you to define rules that specify the source IP address, source port, destination address, destination port, priority, and traffic action (Allow or Deny). The rules can be applied to inbound and outbound traffic independently.

Traditionally, a firewall rule is applied to a port on a router that is connected to a switch. It affects all traffic flowing inbound and outbound to the switch, but it does not affect any traffic within the switch. A Network Security Group rule that is applied to a subnet is more like a firewall rule that is applied at the switch and affects inbound and outbound traffic on every port in the switch. Any virtual machine connected to the switch port would be affected by the Network Security Group rule applied to the subnet.



For example, if a Network Security Group is created and a Network Security Group rule is defined that denies inbound Remote Desktop Protocol (RDP) traffic for all addresses over port 3389, no virtual machine outside the subnet can connect via RDP to a virtual machine that is connected to the subnet, and no virtual machine connected to the subnet can connect via RDP to any other connected virtual machine.

Description	Priority	Source Address	Source Port	Destination Address	Destination Port	Protocol	Action
-------------	----------	----------------	-------------	---------------------	------------------	----------	--------

Deny inbound RDP	1010	*	*	*	3389	TCP	Deny
Allow inbound for subnet	1000	192.168.100.0/24	*	*	3389	TCP	Allow

Network Security Groups can also be applied to the virtual machine or to the network adapter of a virtual machine. This allows greater flexibility in how traffic is filtered. Using network security is strongly recommended for all kind of organizations.

5.2.2 Forced tunneling

Forced tunneling allows you to specify the default route for one or more virtual networks to be the on-premises VPN or ExpressRoute gateway. This results in any packet that is transmitted from a virtual machine connected to the virtual network that is not destined to another IP address within the scope of the virtual network to be sent to that default gateway.

When using forced tunneling, any outbound packet that is attempting to go to an Internet address will be routed to the default gateway and not to the Azure Internet interface. For a virtual machine that has a public endpoint defined that allows inbound traffic, a packet from the Internet will be able to enter the virtual machine on the defined port. A response might be sent, but the reply will not go back out the public endpoint to the Internet. Rather, it will be routed to the default gateway. If the default gateway does not have a route path to the Internet, the packets will be dropped, effectively blocking any Internet access.

Please be aware that forced tunneling has different implementation requirements and scope depending on the type of Azure connectivity of the virtual network. A virtual network that is connected over a S2S VPN connection requires forced tunneling to be defined and configured on a per virtual network basis. A virtual network that is connected over an ExpressRoute connection requires forced tunneling to be defined at the ExpressRoute circuit, and this affects all virtual networks that are connected to that circuit.

5.2.3 Virtual Appliances

Virtual Appliances are third-party-based virtual machine solutions that can be selected from the Azure Gallery or Marketplace to provide services like network firewall, application firewall and proxy, load balancing, and logging. Appliances are licensed by:

- Using a license key that you already own.

- Including the licensing cost into the hourly cost of the appliance.



Appliances are available in single network adapter or multiple network adapter configurations depending on the type of appliance and the required capabilities. The following table lists virtual appliances types and when to use them:

Virtual Appliance Type	Description	When to Use
Network firewall	Virtual appliance that leverages a virtual machine with a multiple network adapter configuration and layer 3 routing support to enable a network firewall between multiple subnets in Azure.	<ul style="list-style-type: none"> • Control outbound traffic flow to the Internet from an application tier • Control inbound traffic flow from the Internet to a UI tier of an application • Control traffic flow between two subnets in Azure • Collect detailed packet captures or network logs of traffic flowing through the appliance
Load balancer	Provides layer 4 or layer 7 load balancing	<ul style="list-style-type: none"> • A load balancer with more features that the Azure Load Balancer is required • Detailed logging is required • SSL termination is required
Security appliance	Intrusion detection appliance	<ul style="list-style-type: none"> • Attempting to create a security stack to manage inbound Internet traffic

		<ul style="list-style-type: none"> Advanced security monitoring and mitigation solution is needed
--	--	--

5.3 Routing of network traffic

Routing of traffic from a virtual machine is accomplished by using implicit system routing via a distributed router that is implemented at the virtual network level. Every packet follows a set of implicit routes that are implemented at the host level. These routes control the flow of traffic within the virtual network to on-premises networks, and to the Internet.

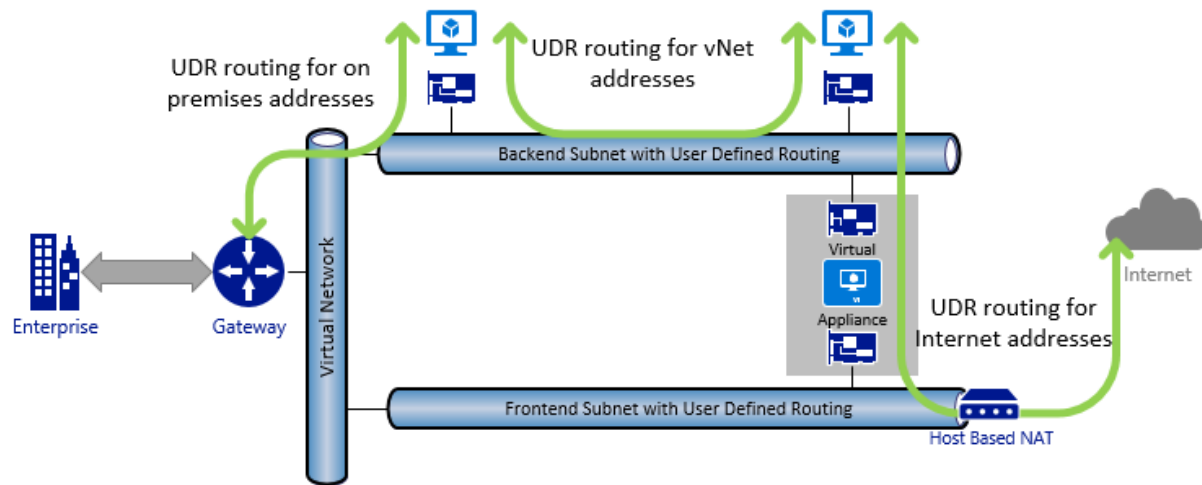
The following rules are applied to the packet in this scenario:

- If the address is within the virtual network address prefix, route to the local virtual network.
- If the address is within the on-premises address prefixes or BGP published routes (BGP or local site network for S2S), route to the gateway.
- If the address is not part of the virtual network, BGP, or local site network routes, route to Internet via NAT.
- If the destination is an Azure datacenter address and ExpressRoute public peering is enabled, it is routed to the gateway because the gateway has the Azure datacenter address via BGP.
- If the destination is an Azure datacenter with S2S or ExpressRoute without public peering enabled, it is routed to the host NAT for the Internet path, but it never leaves the datacenter

User-defined routing allows you to configure and assign routes that override the default implicit system routes, ExpressRoute BGP advertised routes, or the local-site network-defined routes for S2S connections. Configuring a user-defined route allows the specification of next-hop definition rules that control traffic flow within a subnet, between subnets, from a subnet through an appliance to another subnet, to the Internet, and to on-premises networks.

When a network firewall virtual appliance is introduced (see previous section), user-defined routing must be configured to control the traffic routing through the appliance. Without user-defined routing, no traffic will flow through the appliance.

The following diagram shows a virtual appliance inserted into the scenario to control traffic routing to the Internet via front-end and back-end subnets in Azure:



The following rules are applied to the packet in this scenario:

- If the user-defined routing is defined with NextHop Local routing, route to a virtual machine in the virtual network, based on address.
- If the user-defined routing is defined with NextHop VPN Gateway routing, route to a machine on-premises, based on address.
- If the user-defined routing is defined with NextHop Appliance routing, route to the virtual appliance, based on address.
- If the user-defined routing is defined with NextHop Internet routing, route to the Internet over the host NAT.

5.4 Managing public and private IP addresses

IP addresses are assigned to Azure resources to communicate with other Azure resources, in the on-premises network, and the Internet. There are two types of IP addresses you can use in Azure: public and private.

- Public IP addresses are used for communication with the Internet, including Azure public-facing services. There are two methods in which an IP address is allocated to a public IP resource: dynamic or reserved. The default allocation method is dynamic, where an IP address is not allocated at the time of its creation. Instead, the public IP address is allocated when you start (or create) the associated resource (like VM or Load Balancer). The IP address is released when you stop (or delete) the resource. To ensure the IP address for the associated resource remains the same, you can set the allocation method explicitly to reserved. In this case an IP address is assigned immediately. It is released only when you delete the resource or change its allocation method to dynamic.

A typical scenario that requires reserved public IP addresses are Internet-facing websites that are hosted on IaaS machines.

- Private IP addresses are used for communication within an Azure virtual network, and the on-premises network when you use a VPN gateway or ExpressRoute circuit to extend the network to Azure. A private IP address is allocated from the address range of the subnet to which the resource is attached. The address range of the subnet itself is a part of the VNet's address range. There are two methods in which a private IP address is allocated: dynamic or static. The default allocation method is dynamic, where the IP address is automatically allocated from the resource's subnet (using DHCP). This IP address can change when you stop and start the resource. You can set the allocation method to static to ensure the IP address remains the same. In this case, you also need to provide a valid IP address that is part of the resource's subnet. A typical scenario for static private IP addresses are Active Directory Domain Controllers hosted on Azure.

5.5 Recommendations for cloud connectivity

Enterprise organizations benefit from taking a methodical approach to optimizing network throughput across your intranet and to the Internet.

Recommendations for cloud connectivity	
Optimize intranet connectivity to your edge network	Over the years, many organizations have optimized intranet connectivity and performance to applications running in on-premises datacenters. With productivity and IT workloads running in the Microsoft cloud, additional investment must ensure high-connectivity availability and that traffic performance between your edge network and your intranet users is optimal.
Optimize throughput at your edge network	As more of your day-to-day productivity traffic travels to the cloud, you should closely examine the set of systems at your edge network to ensure that they are current, provide high availability, and have sufficient capacity to meet peak loads.
For a high SLA use ExpressRoute	Although you can utilize your current Internet connection from your edge network, traffic to and from Microsoft cloud services must share the pipe with other intranet traffic going to the Internet. In addition, your traffic to Microsoft cloud services is subject to Internet traffic congestion. For a high SLA and the best performance, use ExpressRoute, a dedicated WAN connection between your network and Azure. ExpressRoute can leverage your existing network provider for a dedicated connection. Resources

	connected by ExpressRoute appear as if they are on your WAN, even for geographically distributed organizations
Analyze your current network	<ul style="list-style-type: none"> • Analyze your client computers and optimize for network hardware, software drivers, protocol settings, and Internet browsers. • Analyze your on-premises network for traffic latency and optimal routing to the Internet edge device. • Analyze the capacity and performance of your Internet edge device and optimize for higher levels of traffic. • Analyze the latency between your Internet edge device (such as your external firewall) and the regional locations of the Microsoft cloud service to which you are connecting. • Analyze the capacity and utilization of your current Internet connection and add capacity if needed. Alternately, add an ExpressRoute connection.
Plan and design networking for Azure	<ul style="list-style-type: none"> • Prepare your intranet for Microsoft cloud services. • Optimize your Internet bandwidth. • Determine the type of VNet (cloud-only or cross-premises). • Determine the address space of the VNet. • Determine the subnets within the VNet and the address spaces assigned to each. • Determine the DNS server configuration and the addresses of the DNS servers to assign to VMs in the VNet. • Determine the load balancing configuration (Internet-facing or internal). • Determine the use of virtual appliances and user-defined routes. • Determine how computers from the Internet will connect to virtual machines. • For multiple VNets, determine the VNet-to-VNet connection topology. • Determine the on-premises connection to the VNet (S2S VPN or ExpressRoute). • Determine the on-premises VPN device or router. • Add routes to make the address space of the VNet reachable. • For ExpressRoute, plan for the new connection with your provider. • Determine the Local Network address space for the Azure gateway. • Configure on-premises DNS servers for DNS replication with DNS servers hosted in Azure. • Determine the use of forced tunneling and user-defined routes.

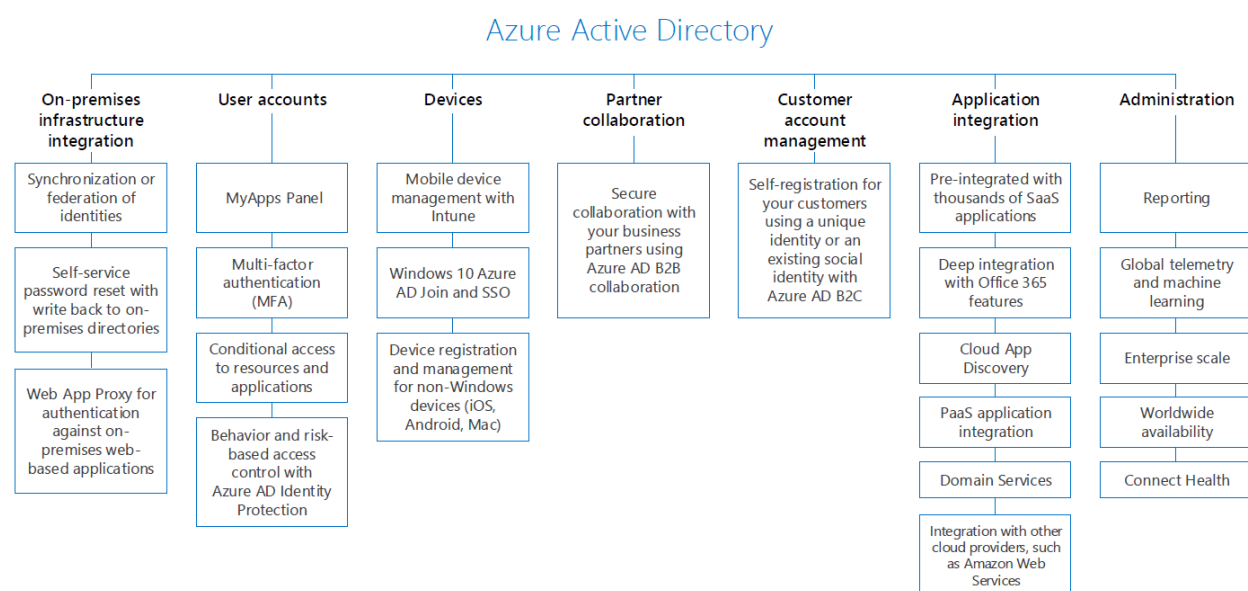
6 Extending Active Directory to Azure

The existence of an Azure Active Directory (Azure AD) tenant is a requirement for any Azure subscription. Therefore, each Azure tenant has at least one Azure AD tenant associated with it. This directory is used for signing in to and accessing Microsoft Azure and other Microsoft cloud services through their corresponding portals, through PowerShell and command line tools, as well as through Graph and Rest APIs.

Azure AD interacts with the cloud in two ways:

- An enabler of the cloud
- A consumer of the cloud

IT professionals will mostly be concerned with Azure AD as an enabler of the cloud because they are often tasked with integrating the enterprise identity and access management platform into the cloud. On the other hand, developers will mostly be concerned with the identity services that Azure AD provides as a consumer of the cloud. Most often, they are looking to understand how their applications can leverage the cloud identity service. The following picture provides an overview of Azure Active Directory.



See [Microsoft Cloud Identity for Enterprise Architects](#).

When a directory is created, the default name of the directory is <something>.onmicrosoft.com. The <something> is chosen by the directory administrator during the creation of the directory. Usually, customers want to use their own domain name, such as contoso.com. This can be achieved by using a custom domain name.

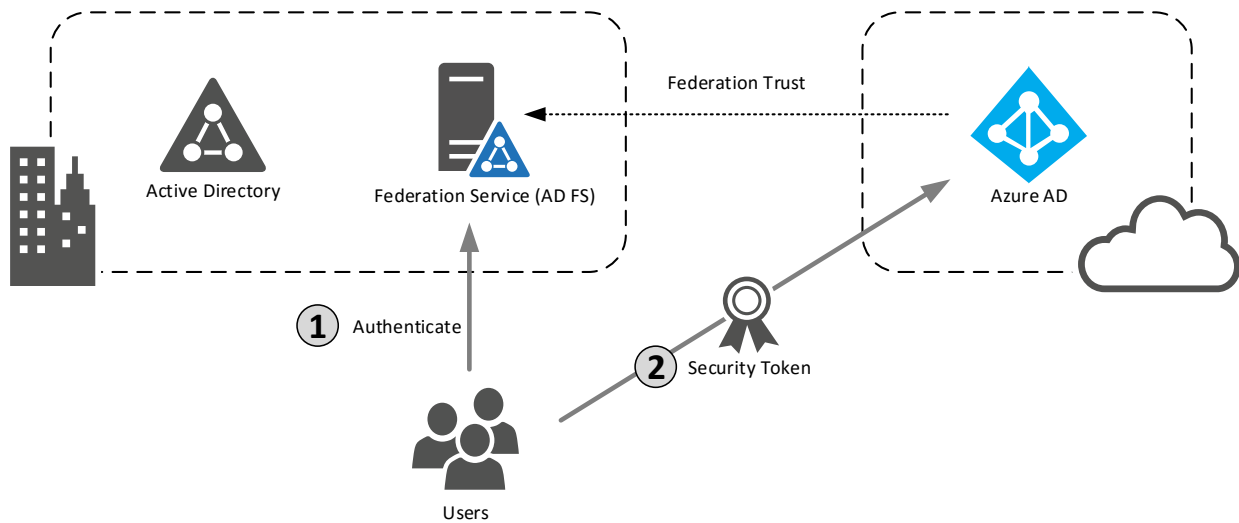
6.1 Synchronizing/federating Active Directory Domain Services with Azure AD

When you create a new Azure AD tenant, the contents of the directory will be managed independently from the on-premises Active Directory forest. This means that when a new user comes in to the organization, an administrator must create an on-premises Active Directory account and an Azure Active Directory account for the employee. Because these two accounts are separate by default, they also may have different user names and passwords, and they need to be managed separately.

However, an organization can use Azure AD Connect to connect the on-premises Active Directory to Azure AD. When this is in place, users who are added or removed from the on-premises Active Directory are automatically added to Azure AD. The user names and passwords are also kept synchronized between the two directories, so end users do not have different credentials for cloud and on-premises systems. With password hash synchronization, the Azure AD Connect service will synchronize one-way SHA256 hashes of Active Directory password hashes into Azure AD. This allows a user that signs into Azure AD to use the same password that is used to sign in to the on-premises Active Directory.

Active Directory Federation Services (AD FS) can be used to add an identity federation trust between on-premises Active Directory and Azure AD. This enables users to have a desktop single sign-on experience when accessing resources that are integrated with Azure AD. With this experience, an end user would sign in to a domain-joined workstation and not be prompted again for a password throughout the entire session, regardless of which applications are used.

When a federation trust is in place, Azure AD defers to the on-premises identity provider to collect the user's credentials and perform the authentication. After authenticating the user, the on-premises identity provider creates a signed security token to serve as proof that the user was successfully authenticated. This security token may also contain data about the user (called claims), which can then be provided to Azure AD for various purposes. The security token is given to Azure AD, which then verifies the signature on the token and uses it to provide access to the applications. The following diagram illustrates this behavior:



Enterprise customers tend to use AD FS to have better control of the authentication of users.

6.2 Working with multiple forests and domains

Most customers do not have simple single-forest Active Directory environments, and dealing with multiple forests can be a challenge when integrating with Azure AD.

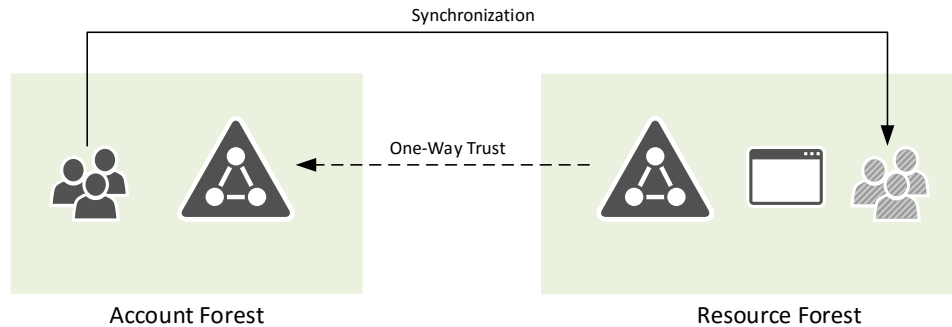
- **Single forest with multiple domains**

Some customers have a single forest environment with multiple domains. Azure AD Connect natively handles this scenario when the following conditions need to be met:

- Users need to exist uniquely across the forest. A user cannot have an active account in more than one domain, because both accounts will be synchronized as separate identities in Azure AD.
- If the domains in the forest use different User Principal Name (UPN) suffixes, each UPN suffix needs to be added to the Azure AD tenant as a custom domain name.

- **Account and resource forest model**

Another common scenario is to have an account and resource forest model. There is a dedicated forest where all of the user identities reside (the account forest) and a dedicated forest for some or all of the applications (the resource forest). A one-way trust (often a forest trust) is in place so that the resource forest trusts the account forest. This relationship is depicted in the following diagram.

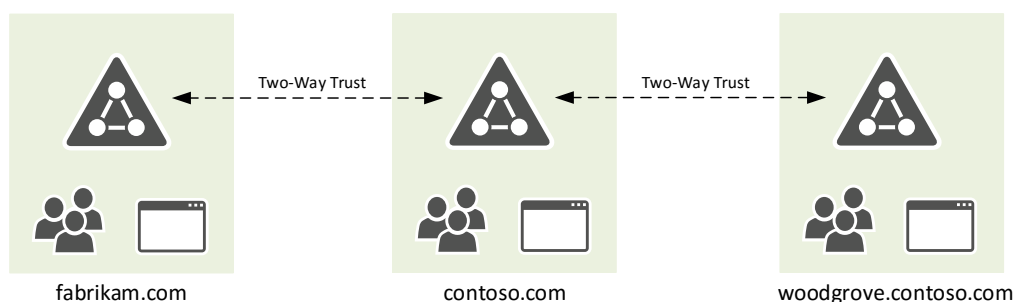


This is most commonly seen with complex Exchange Server deployments. Often, there needs to be a representation of the user in the resource forest's Active Directory for the application to use. This is sometimes referred to this as a *shadow account*. In most cases, it's a duplicate of the user's account from the Account forest, but it is put into a disabled state. Thereby, users are prevented from signing in to it.

Azure AD Connect natively handles this scenario. If the resource forest contains data that needs to be added to Azure AD (such as mailbox information for an Exchange user), the synchronization engine detects the presence of disabled accounts with linked mailboxes. The appropriate data is then contributed to the Azure AD user account.

- **Multiple forests with unique users**

In this scenario, there are multiple independent forests in the environment, which may or may not have Active Directory trust relationships between them. This situation is encountered in highly segmented organizations or companies that acquired other companies via mergers and acquisitions. The following diagram depicts what this architecture might look like.



Users in this scenario have only a single account in one of the forests (they do not have multiple user accounts across forests). Because of this, there is no need for the synchronization tool to match a user to multiple accounts.

However, one decision that needs to be made is whether the accounts will be migrated into a single forest at some point. This is an important thing to consider, because it will

determine whether you can use the objectGUID of the user accounts as the source anchor (which is used to match the Active Directory accounts to the Azure AD accounts).

If the users will be migrated to a single forest at some point, you'll need to use a different source anchor, such as the user's email address or UPN. The reason is that the objectGUID can't be migrated with the user. After migration, there would be multiple accounts in Azure AD for migrated users—one for the old forest and another for the new forest.

- **Multiple forests with duplicate users**

This scenario is the same as the previous scenario (multiple forests with unique users) with the exception that a single user has multiple user accounts in different forests in the environment. These accounts are either:

- Enabled (users likely have a password and sign in to these accounts)
- Disabled (a shadow account is used to store attributes for an application, such as Exchange).

Even though there are multiple user accounts in the organization, there should be only a single account for the user in Azure AD. To enable this, the synchronization service needs to be able to match user accounts across the forests to a single person. For this to happen, the accounts in each forest need to have an attribute that contains the same, unique value for a user.

When enabling a federated identity relationship between Azure AD and an on-premises identity provider, an entire domain name in Azure AD is converted from a standard domain to a federated domain. This impacts all of the users that have UPNs under the domain name. You cannot have a mix of federated and nonfederated users in a domain name. Any subdomains under a domain namespace will have the same configuration as the parent domain.

After the domain name is converted to federated, all users who attempt to sign in to Azure AD with a UPN from the converted domain (or one of its child domains) will be redirected to the on-premises identity provider for authentication. If the user does not have a valid account in the on-premises identity provider, the user will not be able to authenticate to Azure AD or to any of the connected applications.

AD FS can support a multiple forest configuration, but only if all the forests have two-way Active Directory trust relationships between them. If there are no forest trusts between the multiple Active Directory Domain Services (AD DS) forests, you must have multiple AD FS deployments (one for each forest that is untrusted). If possible, we recommend that customers have trusts between their multiple Active Directory forests, so that only a single AD FS farm is needed for Azure AD.

6.3 Multi-Factor Authentication

Multi-Factor Authentication (MFA) is a method of authentication that requires the use of more than one verification method and adds a critical second layer of security to user sign-ins and transactions. It works by requiring any two or more of the following verification methods:

- Something you know (typically a password)
- Something you have (a trusted device that is not easily duplicated, like a phone)
- Something you are (biometrics)

Azure Multi-Factor Authentication is a method of verifying who you are that requires the use of more than just a username and password. It provides a second layer of security to user sign-ins and transactions. Azure Multi-Factor Authentication helps safeguard access to data and applications while meeting user demand for a simple sign-in process. It delivers strong authentication via a range of easy verification options—phone call, text message, or mobile app notification or verification code and third-party OAuth tokens.

The security of Multi-Factor Authentication lies in its layered approach. Compromising multiple authentication factors presents a significant challenge for attackers. Even if an attacker manages to learn the user's password, it is useless without also having possession of the trusted device. Should the user lose the device, the person who finds it won't be able to use it unless he or she also knows the user's password. Azure Multi-Factor Authentication is available in three different versions:

- **Multi-Factor Authentication for Office 365**
This version works exclusively with Office 365 applications and is managed from the Office 365 portal. So administrators can now help secure their Office 365 resources by using Multi-Factor Authentication. This version comes with an Office 365 subscription.
- **Multi-Factor Authentication for Azure Administrators**
The same subset of Multi-Factor Authentication capabilities for Office 365 will be available at no cost to all Azure administrators. Every administrative account of an Azure subscription can get additional protection by enabling this core Multi-Factor Authentication functionality. So administrators who want to access the Azure portal to create a VM, a website, manage storage, mobile services or any other Azure Service can add Multi-Factor Authentication to their administrator account. It is best practice that all Azure administrator accounts should be configured for MFA.
- **Azure Multi-Factor Authentication**
Azure Multi-Factor Authentication offers the richest set of capabilities. It provides additional configuration options via the Azure Management portal, advanced reporting,

and support for a range of on-premises and cloud applications. Azure Multi-Factor Authentication comes as part of Azure Active Directory Premium and is also available as a stand-alone service with per user and per authentication billing (see <https://azure.microsoft.com/en-us/pricing/details/multi-factor-authentication>).

6.4 Hosting Active Directory domain services

When working with virtual machines in an Infrastructure-as-a-Service (IaaS) environment, the virtual machines most often need to be joined to an Active Directory domain. This is required so the operating system can be properly managed and the software running on the virtual machines can function properly. Many customers who move virtual machines to Azure have come to the conclusion that extending Active Directory into Azure IaaS is a recommended course of action.

One of the questions we are often asked is whether a customer should deploy Active Directory domain controllers into IaaS. The alternative option is to keep them on-premises and provide a VPN connection. There are various considerations to be made when answering this question.

Whether the domain controllers are on-premises or deployed in Azure, there needs to be connectivity between the Azure virtual network and the on-premises network. If you want to keep domain controllers on-premises, you need an ExpressRoute connection or a Site-to-Site VPN connection to Azure. Every time a virtual machine in Azure needs to access a domain controller, it will traverse this connection over the WAN. Depending on the stability, performance, and latency of the connection, this may cause issues.

Most customers will strongly consider placing domain controllers in Azure because they will want the applications they place in Azure IaaS to have reliable and low latency access to the domain controllers. Domain controllers are highly sensitive roles. If someone compromises a domain controller, they can gain access to virtually everything in a customer's environment. The best way to handle the situation is to use the Tier 0 subscription to host the domain controllers as described in section 4.2. It allows you to manage who has explicit control over the domain controllers and their security equivalents in Tier 0.

When deploying domain controllers in Azure, there are some specific things to consider for your Active Directory design.

- It is generally recommended to consider the Azure datacenter as a separate Active Directory site, because it will have its own IP address space and routing considerations. For many applications and services, it is preferable to have a domain controller available within the site, and it's typically preferable to have a local connection instead of traversing the WAN.

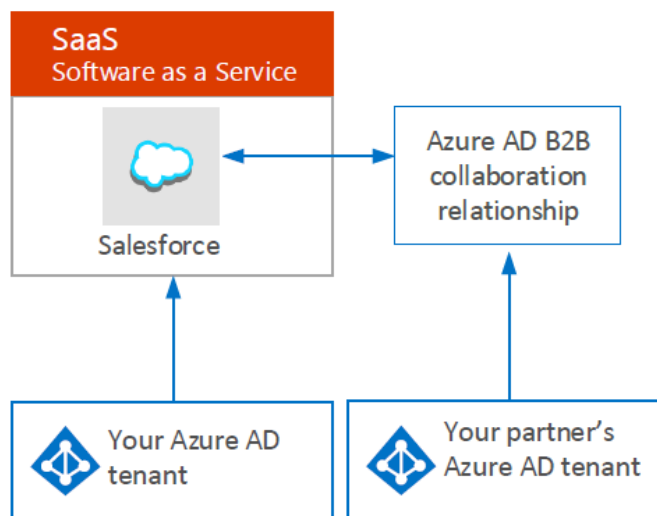
- More important is the need to consider what happens to a virtual machine in Azure if it can't reach a domain controller. The standard recommendation is to place two domain controllers within an availability set in each Azure region where virtual machines reside. It is important to note that a resource domain or forest is not recommended given the additional overhead, and these do not represent an effective security boundary.
- In addition, there should be an Active Directory site created for each Azure region, and all of the virtual networks in that region should be associated with that site. Standard guidance applies for the definition of Active Directory site links.
- In modern Active Directory deployment there's little reason to not make every domain controller a Global Catalog server. The standard guidance for Global Catalogs also applies to domain controllers in Azure. As a recommended practice, make all of the domain controllers in Azure Global Catalog servers.
- DNS is instrumental to the operation of Active Directory. There should always be DNS servers located alongside the domain controllers, and most of the time we recommend that DNS be Active Directory-integrated. This does not change with Azure. The domain controllers in Azure should run the DNS Server service, if possible. If you are not using DNS in Windows, there should be a DNS appliance in Azure for the domain controllers to use. Otherwise, a VPN outage will render DNS unavailable and prevent the domain controllers in Azure from operating correctly. DNS Servers need to be registered in the Azure virtual networks.
- Azure provides a default DNS service to virtual machines if you don't specify a DNS server. The Azure name resolution services do not support the complex name resolution needs of Active Directory, so do not attempt to use Azure DNS servers on domain controllers.
- Organizational units (OUs) in an Active Directory design are important for operational and security management of Azure assets, especially when extending existing on-premises forests into the Azure cloud. OUs, security groups, and Group Policy Objects (GPOs) provide key administrative controls that can provide containment boundaries within a security zone.

6.5 Using additional Azure Active Directory elements

6.5.1 Azure AD B2B Collaboration

[Azure AD B2B Collaboration](#) enables secure collaboration between business-to-business partners. These new capabilities make it easy for organizations to create advanced trust

relationships between Azure AD tenants so they can easily share business applications across companies without the hassle of managing additional directories or the overhead of managing partner identities. With 6 million organizations already using Azure AD, chances are good that your partner organization already has an Azure AD tenant, so you can start collaborating instantly. But even if it doesn't, Azure AD's B2B capabilities make it easy for you to send it an automated invitation that will get it up and running with Azure AD in a matter of minutes.



6.5.2 Azure AD B2C Collaboration

[Azure Active Directory B2C](#) is a highly available, global, identity management service for consumer-facing applications that scales to hundreds of millions of identities. It can be easily integrated across mobile and web platforms. Your consumers can log on to all your applications through fully customizable experiences by using their existing social accounts or by creating new credentials.

6.5.3 Azure AD Domain Services

Azure AD Domain Services provides managed cloud-based domain services such as domain join, group policy, LDAP, and Kerberos/NTLM authentication in Azure IaaS that are fully compatible with Windows Server AD. You can join Azure virtual machines to this domain without the need to deploy domain controllers. Because Azure AD Domain Services is part of your existing Azure AD tenant, users can login using the same credentials they use for Azure AD. This managed domain is a standalone domain and is not an extension of an organization's on-premises domain or forest infrastructure. However, all user accounts, group memberships, and credentials from the on-premises directory are available in this managed domain.

6.5.4 Azure Application Proxy

[Microsoft Azure Active Directory Application Proxy](#) lets you publish applications, such as SharePoint sites, Outlook Web Access, and IIS-based apps, inside your private network and provides secure access to users outside your network. Employees can log into your apps from home on their own devices and authenticate through this cloud-based proxy. By using Azure AD Proxy you can protect on-premises applications with the same requirements as other cloud-based applications with MFA, device requirements, and other conditional access requirements. You also benefit from the built-in security, usage, and administration reports. Application Proxy works by installing a slim Windows service called a Connector inside your network. The Connector maintains an outbound connection from within your network to the proxy service. When users access a published application, the proxy uses this connection to provide access to the application.

6.6 Recommendations for using Azure Active Directory

Recommendations for using Azure Active Directory	
Follow best practices for setting up Active Directory Domain Services in Azure	<ul style="list-style-type: none"> • Create a unique Active Directory site object for each Azure region where virtual machines reside, and associate all of the virtual networks in that region with the Active Directory site. • Place two domain controllers within an availability set in all Azure regions where virtual machines reside. • Make all domain controllers in Azure Global Catalog servers. • Make sure that domain controllers are pointing to a DNS server in Windows that hosts the Active Directory zones, rather than the default DNS servers in Azure. • Do not set a static IP address on the network adapter in the operating system for virtual domain controllers in Azure. Doing so will isolate the virtual machines and prevent them from communicating on the virtual network. • To give a domain controller the IP address that you want and prevent it from changing if the virtual machine is de-provisioned, provide the virtual machine with a static virtual network IP address. • Make sure that you place the Active Directory database and SYSVOL on a data disk. If you use the operating system disk or a temporary disk, the database may get corrupted or purged during an outage.

Enable password hash synchronization	Enable password hash synchronization so that the Azure AD password for users is the same as the on-premises Active Directory password. Even if all of a customer's users are signing in to Azure AD with AD FS, it is recommended to enable password synchronization. Doing so provides a good fallback method for user authentication if AD FS goes offline
Prepare your Active Directory	<ul style="list-style-type: none"> • If users from the additional forests will be migrated into a single forest in the future, you must choose something other than the objectGUID as the source anchor attribute (such as the mail attribute). • If a single person has multiple user accounts in different forests, you must choose a common attribute to match the accounts together. • If user certificates use the UPN in the Subject Name field, the certificates need to be reissued during the UPN rationalization.
Plan your AD FS deployment properly	<ul style="list-style-type: none"> • Use a single identity provider for the organization, if possible. Otherwise, you'll have to manage multiple instances of the Identity Federation Service on-premises. • If possible, we recommend that you have trusts between each Active Directory forest and use a single AD FS instance with Azure AD. This simplifies the architecture and prevents you from having to manage multiple AD FS farms. • Use Web Application Proxy servers in the AD FS deployment for Azure AD. As a general rule, we recommend starting with an equal number of Web Application Proxy servers and AD FS servers. • If AD FS is used in a multiple forest configuration with trusts between the Active Directory forests, the UPN suffixes for each domain must be unique. • If using AD FS in a multiple forest configuration with no forest trusts between Active Directory forests, you must have multiple deployments of AD FS (one for each untrusted forest).
Use Multi-Factor Authentication	Multi-Factor Authentication is an optional service that increases the security of Azure AD.
Use self-service password reset	We recommend that you create an end-user communication plan to provide the users with the details about how to register for self-service password reset, reset their password, and know what to expect.

7 Operating Azure IaaS Services

As enterprise IT teams broaden their server deployments on-premises, in the cloud, or in a hybrid-model, these bring forth many management challenges. To overcome these IT management challenges, organizations are utilizing today disjointed solutions for individual management needs. However, customers can now take advantage of cloud scale infrastructure and increase ease of deployment through a unified “IT Management as a Service,” which is called Microsoft Operations Management Suite (OMS). Management capabilities such as monitoring, backup, automation, and so forth are delivered as a service from the cloud that connects all of the servers in all environments (on-premises, Azure, and other clouds such as AWS) and allows IT staff to centrally manage operations. OMS consists of the following 4 modules:

- Log Analytics → Gain visibility across your Hybrid Enterprise Cloud
- Automation → Orchestrate complex and repetitive operations
- Availability → Increase data protection and application availability
- Security → Help secure your workloads, servers, and users

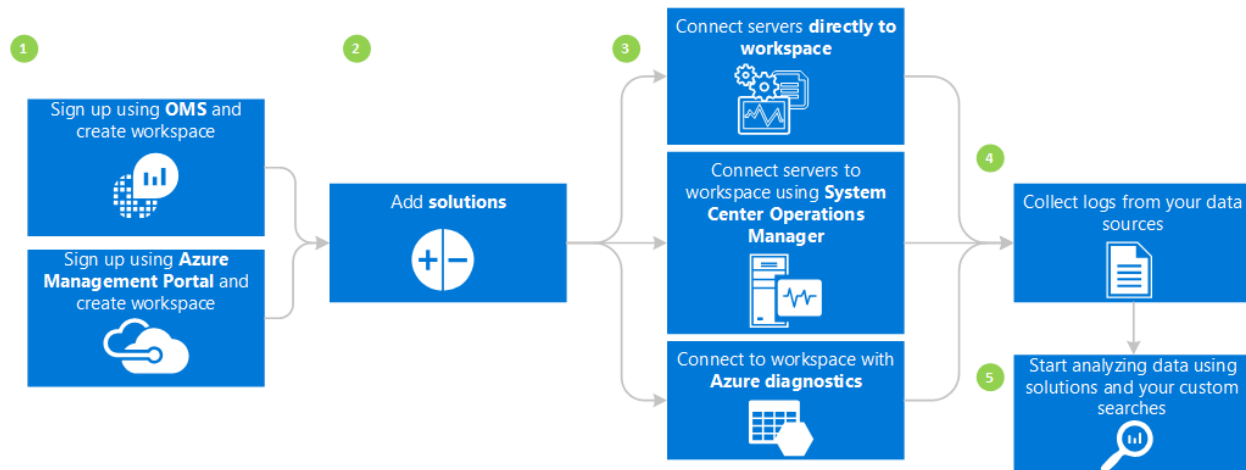
It is not intended that OMS is replacing Microsoft System Center. It extends the capabilities of Systems Center to deliver a full hybrid management experience. This guide is focusing on OMS.

7.1 Gaining operational insights

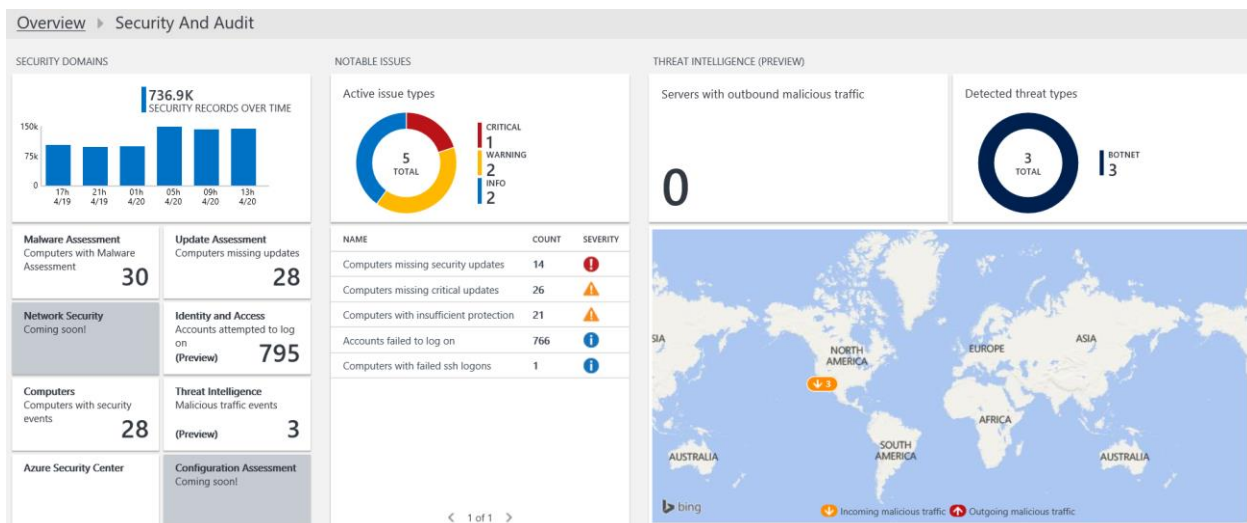
Today traditional IT is usually using multiple different tools for platform and application monitoring, network monitoring, and Security Analysis. Extending those tools to the cloud is challenging in various aspects: connectivity, agility, and data volume. Furthermore, it is more and more necessary to combine and analyze information from various sources to gain operational insights. With OMS Log Analytics, organizations can collect, store, and analyze log data from virtually any Windows Server and Linux source and get unparalleled insights across their datacenters and clouds, including Azure and AWS.

7.1.1 Getting started with Log Analytics

Log Analytics is new to most enterprises and thus it might be helpful to get a quick introduction. To get started with Log Analytics you need to perform the following steps:




The first step is the deployment of an OMS workspace in the Azure Portal. Thereby you can choose from various editions defining the volume and retention of your data. Secondly you are selecting the solutions that you want to use. Solutions are a collection of logic, visualization, and data acquisition rules that address key customer challenges. They allow deeper insights to help investigate and resolve operational issues faster, collect and correlate various types of machine data, and help you be proactive with activities such as Change Tracking, Patch status reporting, and security auditing.



The following picture shows the available solutions at the time of this writing. The gallery is continuously extended with additional solutions. Those can be added at any point in time.

AD Assessment Free Assess the risk and health of Active Directory environments.	Alert Management Free Manage your Operations Manager alerts across your servers.	Automation Free Automate time consuming and frequently repeated tasks in the cloud and on-premises.	Backup Free Manage Azure IaaS VM backup and Windows Server backup status for your backup vault.	Upgrade Analytics Private Preview Coming Soon Mitigate and learn about compatibility issues and upgrade to a later version of Windows.	Key Vault Coming Soon Understand your Key Vault usage through Analysis of Key Vault logs	Office 365 Coming Soon Get full visibility into your Office 365 user activities, perform forensics as well as audit and compliance.	Service Fabric Coming Soon Identify and troubleshoot issues across your Service Fabric cluster	SQL Assessment Free Assess the risk and health of SQL Server environments.	System Update Assessment Free Identify missing system updates across your servers.
AD Replication Status Free Identify Active Directory replication issues in your environment.	Malware Assessment Free View status of antivirus and antimalware scans across your servers.	Azure Networking Analytics Coming Soon Gain insight into your Azure Network data	Change Tracking Free Track configuration changes across your servers.	Containers Coming Soon See Docker container performance metrics and logs from containers across your public or private cloud environments.	Network Performance Monitor Coming Soon Offers near real time monitoring of network performance parameters like loss and latency.	Security and Audit Free Provides the ability to explore security related data and helps identify security breaches.	Azure Site Recovery Free Monitor virtual machine replication status for your Azure Site Recovery Vault.	Surface Hub Free Provides the ability to monitor Microsoft Surface Hub devices.	Wire Data Coming Soon Provides the ability to explore wire data and helps identify network related issues.

Afterward you are connecting your servers to Log Analytics by deploying the Windows or Linux agent on your machines. For Azure virtual machines it is an automated process; for virtual machines in other clouds or on-premises it could be done by downloading and installing the agent or by using automation tools therefore. The connectivity between the agent and Log Analytics is established by using outbound Internet connectivity with or without a proxy. Thus there is no need to open any on-premises firewalls for establishing inbound connectivity.



For those machines that don't have Internet connectivity you can use the Log Analytics Forwarder (Gateway), which is currently in preview. The OMS Log Analytics Forwarder enables you to send data to a central server on your premises, which has access to the Internet and acts as an http forward proxy. This allows you to collect log files from any machine in your network without having the need for Internet access. In addition, you can grab logs that are stored from other Azure services in Azure storage accounts and include them into the repository. Next you define the data that you want to collect from your virtual machines. You can choose to add any Windows Event log, Windows Performance Counters, Linux Performance Counters, IIS

logs, Custom fields, Custom logs, and Syslog. All of those log files are transferred with low latency and stored in a central repository for further analysis. User Accounts can be added to OMS Log analytics according to the RBAC model to define granular access for various users and user groups.

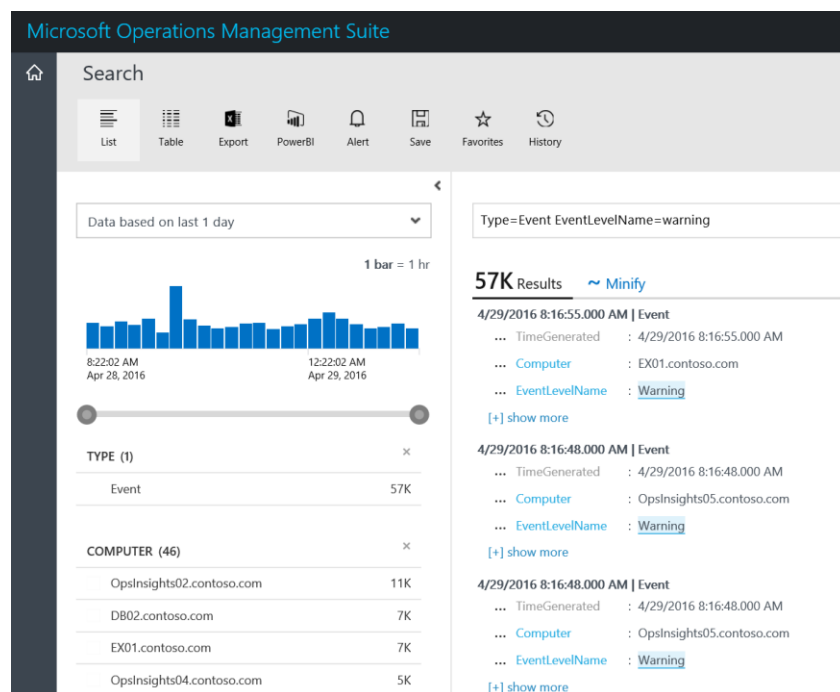
7.1.2 Creating log searches and raising of alerts

At the core of OMS is the log search feature, which allows you to combine and correlate any machine data from multiple sources within a hybrid environment. Solutions are also powered by log search to bring you metrics pivoted around a particular problem area. Throughout the OMS console, you can click tiles or drill in to other items to view details about the item by using log search.

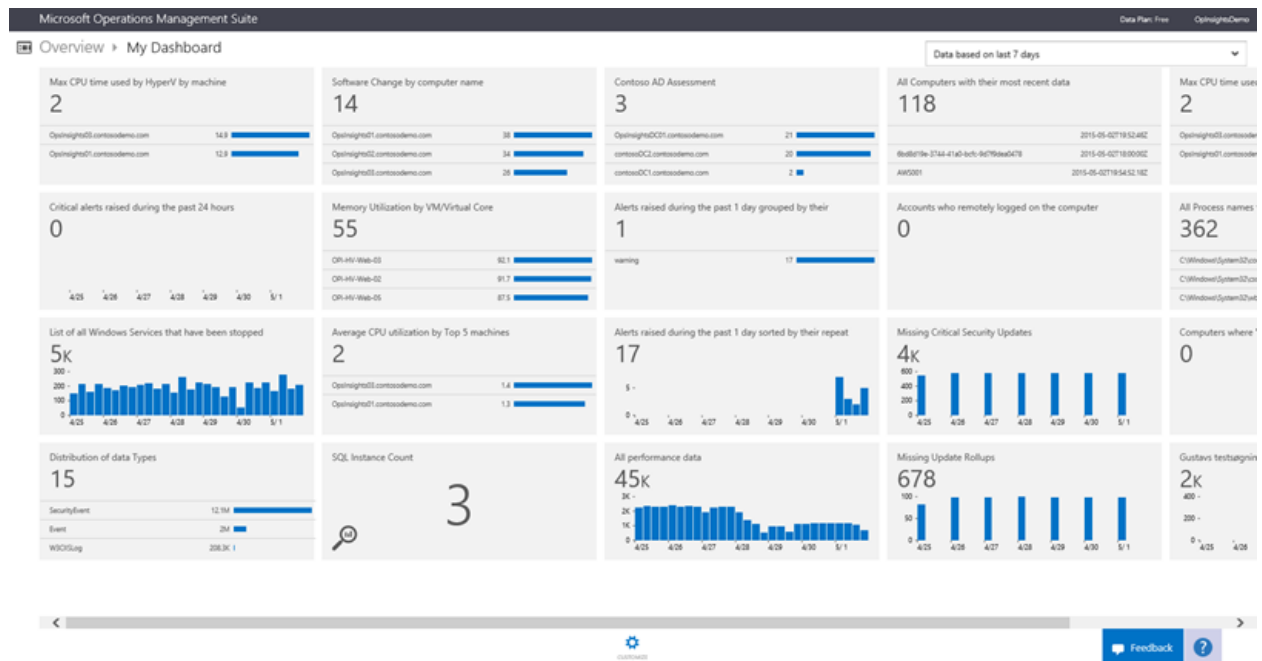
Log search provides a rich set of functions such as:

- Basic and advanced filters
- Measure commands to apply statistical functions to your data and aggregate results
- Max and Min, average, and sum functions
- Searches and Sub-searches

The following screenshot shows a custom query and its results. A complete tutorial to create your own queries can be found here: <https://technet.microsoft.com/en-us/library/mt484120.aspx>.



Queries can be saved and added to favorites and their history is kept in OMS. Result sets can be displayed in different styles and exported in different formats. Furthermore, queries are the foundation to create custom dashboards to highlight those data that are relevant for your business.



One of the most important capabilities of Log Analytics is Alert Rules. Alert rules are based on log searches that you create. They can automatically inform you by sending an email notification, and OMS can remediate issues with Automation runbooks. Alert rules in OMS are based on saved log search queries, the frequency that the alert rule runs, the time range of the alert rule, and a condition based on the number of results for the query. Since an alert rule is based on a search query, you can use some saved search queries before you create an alert rule, or you can use an active search query to get started with a new alert rule. For example, you can create an alert rule to notify you when more than 10 warning events were generated over the past hour on a specific virtual machine and to run the alert rule every 30 minutes. The usefulness of OMS alerts is that after you create them, you don't have to manually check when those important events occur—instead, OMS continuously looks for those important events and can immediately inform you when they occur. And, OMS can run any runbook job from your Automation account to remediate the problem in an automated fashion wherever possible.

The email notification of an alert contains all relevant information that is required to create an Incident in your Service Management tool of choice that you are using for your Information Technology Infrastructure Library (ITIL) operations.

7.1.3 Securing data

Microsoft is committed to protecting your privacy and securing your data, while delivering software and services that help you manage the IT infrastructure of your organization. We recognize that when you entrust your data to others, that trust requires rigorous security. Microsoft adheres to strict compliance and security guidelines—from coding to operating a service. Further details are outlined in section 3.

The OMS service manages your cloud-based data securely by using the following methods:

- **Data segregation**
Customer data is kept logically separate on each component throughout the OMS service. All data is tagged per organization. This tagging persists throughout the data lifecycle, and it is enforced at each layer of the service. Each customer has a dedicated Azure blob that houses the long-term data.
- **Data retention**
Aggregated metrics for some of the solutions such as Capacity Management are stored in a SQL Database hosted by Microsoft Azure. This data is stored for 390 days. Indexed log search data is stored and retained according to the pricing plan.
- **Physical security**
The OMS service is manned by Microsoft personnel, and all activities are logged and can be audited. The OMS service runs completely in Azure and complies with the Azure common engineering criteria.
- **Compliance and certifications**
The OMS software development and service team is actively working with the Microsoft Legal and Compliance teams and other industry partners to acquire a variety of certifications.

OMS Log Analytics currently meet the following security standards:

- Windows Common Engineering Criteria
- Microsoft Trustworthy Computing Certification
- ISO/IEC 27001 compliant
- Service Organization Controls (SOC) 1 Type 1 and SOC 2 Type 1 compliant

7.2 Backing up and restoring data

Backing up and restoring data are key for any production and most nonproduction workloads. The relevant scenarios are:


- Backup and restore a virtual machine
- Backup and restore files and folders
- Backup and restore application data

Instead of investing in a new or extending an on-premises backup solution, which needs to be engineered, operated, and maintained, customers can take advantage of Azure Backup, a cloud scale backup infrastructure that is easy to use and managed by Microsoft. Some of the key benefits of a cloud-based backup are:

- No capital expenditure is needed for on-premises storage devices
- Pay-as-you-use consumption model
- Unlimited scaling
- Multiple storage options
- Data encryption
- Application consistent backup
- Long-term retention

7.2.1 Azure virtual machines

Backing up and restoring business-critical data is complicated by the fact that it needs to be backed up while the applications that produce the data are running. To address this, Azure Backup provides application-consistent virtual machine backups for Microsoft workloads by using VSS to ensure that data is written correctly to storage. For Linux virtual machines, file-consistent backups are possible, since Linux does not have an equivalent platform to VSS.



The foundation of every Azure Backup is the Recovery Services Vault. The vault could be based on Geo-Redundant Storage (GRS) or Locally-Redundant Storage (LRS). GRS provides three copies of your backup data in the primary region and three more copies in another region to be protected from regional disasters. A best practice is to use Geo-Redundant-Storage, especially for any production workloads. The storage type needs to be configured before you start protecting machines.

Azure Backup discovers all machines in a subscription and allows you to add new machines to a backup schedule. The backup schedule is defined by backup policies. There is a default policy, representing a typical schedule along with the option to create multiple different custom backup policies. The backup policy defines the backup frequency, as well as the daily, weekly, monthly, and yearly retention period.

* Policy Name ⓘ

BACKUP FREQUENCY

Local Time (UTC+02:00)

RETENTION RANGE

☒ DAILY ☒ WEEKLY ☒ MONTHLY ☒ YEARLY

DAILY RETENTION

* At For Day(s)

WEEKLY RETENTION

* On * At For Week(s)

MONTHLY RETENTION

* On * * At For Month(s)

YEARLY RETENTION

* In * On * * At For Year(s)

The most time-consuming operation in backup is the copying of data from the primary storage to the backup storage, and the time taken is dependent on a host of factors like network latency, available IOPS on the primary storage account, and available IOPS on the backup storage account. Azure Backup implements an optimized blob copy that ensures constant, predictable IO and backup times. Azure Backup does additional processing to determine the

incremental changes between the last recovery point and the current VM state. By transferring and storing only the incremental changes, Azure Backup is highly storage efficient.

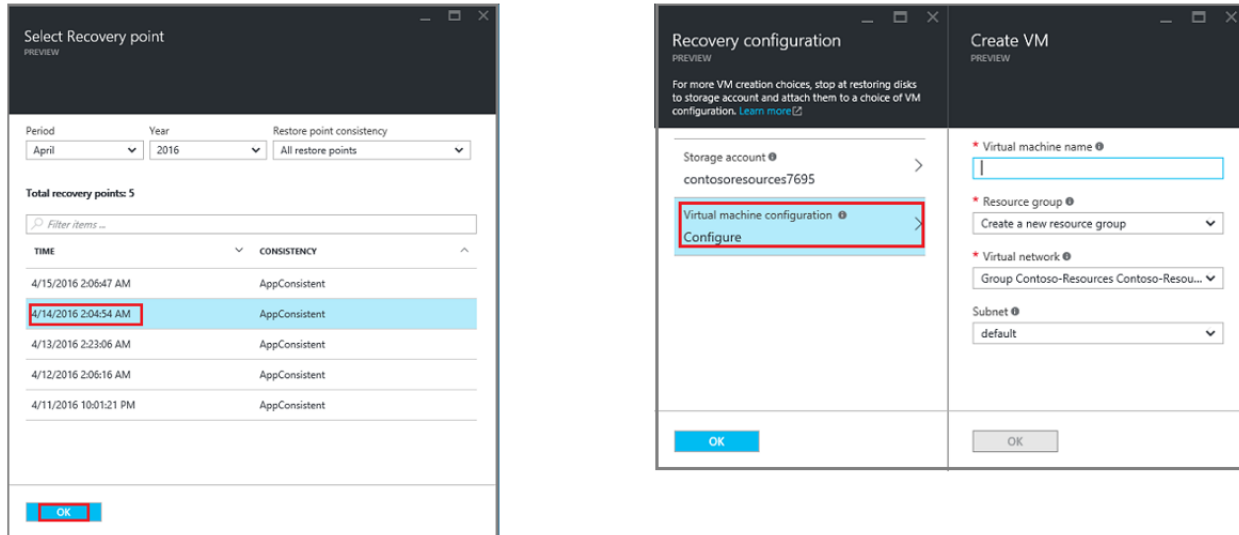
For enterprises looking to encrypt their VM data in Azure, the solution is to use BitLocker on Windows or dmccrypt on Linux machines. Both of these are volume-level encryption solutions. The entire encryption of data happens transparently and seamlessly in the VM layer. Thus the data written to the page blobs attached to the VM is encrypted data. When Azure Backup takes a snapshot of the VM's disks and transfers data, it copies the encrypted data present on the page blobs.

Azure backup provides auditing capabilities for backup operations triggered by the customer, making it easy to see exactly what management operations were performed on the backup vault. Operations logs enable great post-mortem and audit support for the backup operations.

The following operations are logged:

- Register and Unregister
- Configure protection
- Backup and Restore
- Stop protection
- Delete backup data
- Add, delete, update policy
- Cancel job

Even more important than backing up your virtual machines is the ability to restore the entire virtual machine. You protect your data with the Backup service by taking snapshots of your data at defined intervals. These snapshots are known as recovery points, and they are stored in recovery services vaults. If or when it is necessary to repair or rebuild a VM, you can restore the VM from any of the saved recovery points. When you restore a recovery point, you return or revert the VM to the state when the recovery point was taken. The restore configuration allows you to specify the name, resource group, network, and subnet for the virtual machine.



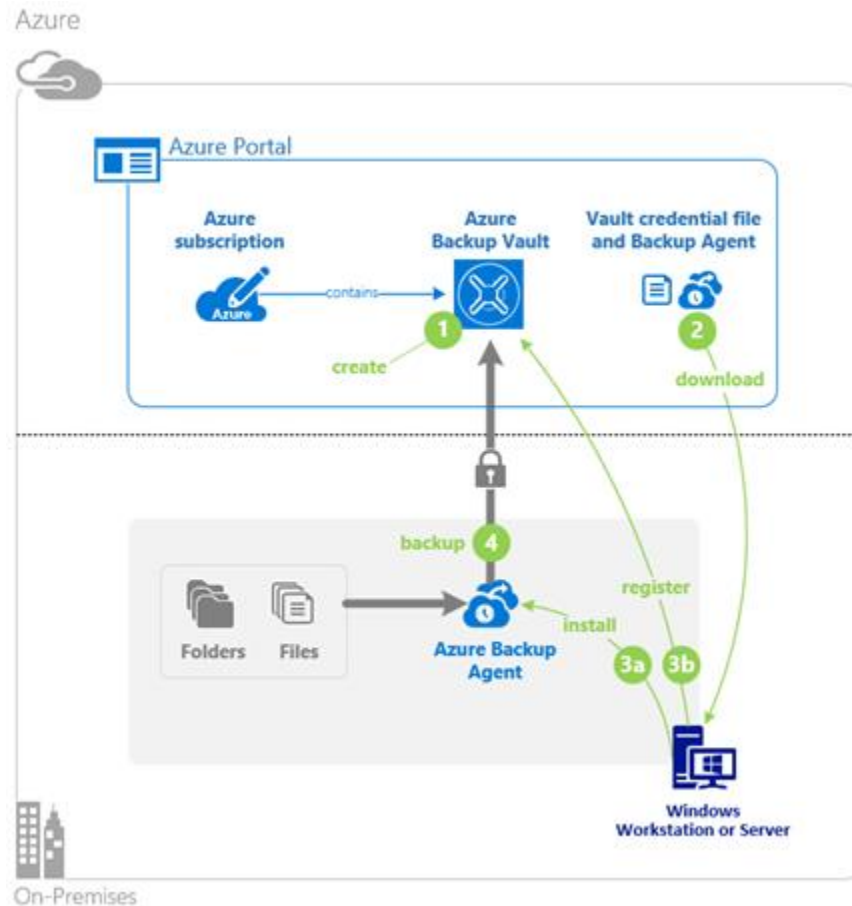
Complex restore scenarios such as VMs under load balancer, VMs with multiple reserved IPs, and VMs with multiple NICs require usage of PowerShell commands as described here:

<https://azure.microsoft.com/en-us/documentation/articles/backup-azure-vms-automation/#restore-an-azure-vm>


7.2.2 Files and folders

This backup option is designed to back up files and folders from any Windows machine. The machine can run in Azure, on-premises, or in any other cloud; it can be physical or virtual. You cannot use this option to back up the system state, or to create a Bare-Metal-Restore (BMR) backup. The Recovery Services Vault could be the one that is mentioned in the previous section, or it could be any other Recovery Services Vault.

Azure Backup for files and folders requires the installation of an Azure Backup Agent on the server, which can be downloaded from the Azure Recovery Services Vault. After installing the agent, it is necessary to connect the server to the Recovery Services Vault by downloading the vault credential files from the Recovery Services Vault. The vault credentials file is used only during the registration workflow and expires after 48 hours. Ensure that the vault credential file is available in a location that can be accessed by the setup application.



The backup agent provides a user interface to schedule multiple backups per day/week, to define retention policies (according to the previous section). The Azure Backup agent is establishing connectivity to the Recovery Services Vault by using https over the Internet with or without a proxy server. Furthermore, it is possible to use ExpressRoute Public peering to establish the connection over a private WAN link. The Backup agent provides network throttling. Throttling controls how network bandwidth is used during data transfer. This control can be helpful if you need to back up data during work hours but do not want the backup process to interfere with other Internet traffic. Throttling applies to backup and restore activities.



A key to overcoming concerns about security in the cloud is encryption. Azure Backup requires a passphrase for encrypting data to be backed up. The agent performs encryption prior to transmission to Azure, and decryption after performing a restore operation back to the local system. Microsoft advises that it does not keep a copy of the passphrase. Azure Backup uses compression to reduce transmission time and storage requirements. Once the compression and encryption is applied, the data in the backup vault is usually 30 to 40 percent smaller. Azure Backup also uses block-level incremental backup methods so that only modified blocks are backed up in subsequent backups of the same set of files and folders.

When restoring data to the same server from which it was backed up, you don't have to supply the passphrase, but when restoring to a different server, you do. Azure Backup can generate a passphrase for you, or you can create your own; in either case, you can change it in the console's Properties page later.

7.2.3 Enterprise applications

Besides the need of backing up virtual machines and files/folders, it is also required to back up enterprise applications in a consistent way and to restore parts of the protected application in a granular way. A lot of Systems Center customers are using Data Protection Manager for enterprise application backup on-premises and in the cloud. Large customers, who have another backup strategy for their on-premises applications, struggle to introduce Data Protection Manager due to the required Systems Center licenses.

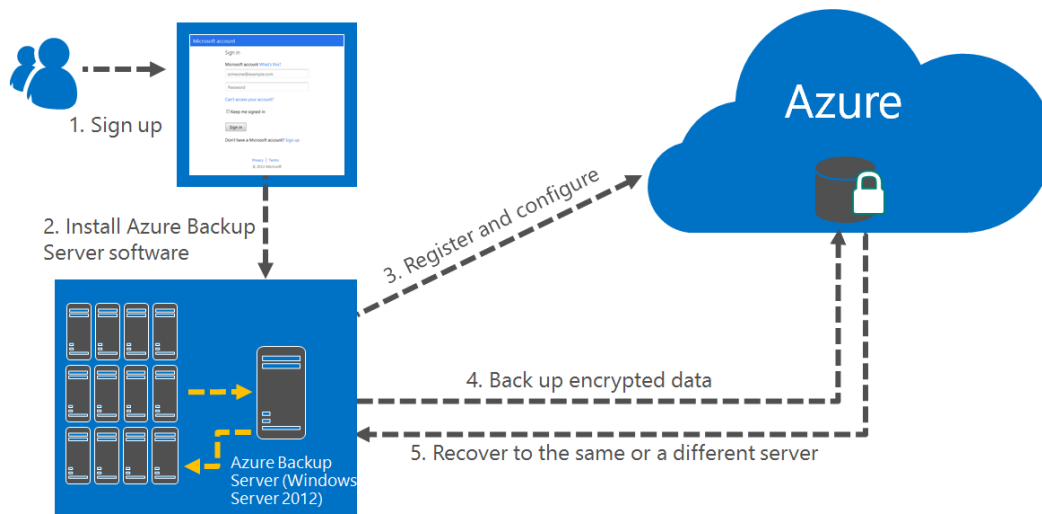
Microsoft Azure Backup server (MABS) is a new tool that provides disk to disk to cloud backup with centralized local management and economic cloud-based offsite storage. It supports to

protect application workloads such as Virtual Machines (on-premises), Microsoft SQL Server, SharePoint Server, Microsoft Exchange, and Windows clients from a single console. Azure Backup Server inherits the functionality of Data Protection Manager (DPM) for workload backup. The main differences between DPM and MABS are:

- DPM offers tape protection, which is not available in MABS.
- DPM can protect one datacenter's DPM installation with a secondary DPM server in another datacenter (and vice versa), which MABS doesn't offer.
- Many DPM servers can be managed in a single, central console in Operations Manager.
- DPM can act as a conduit for Azure Site Recovery services with Hyper-V replica, whereas MABS only does backup.
- MABS requires an Azure Backup Subscription.

This document is focusing on using MABS for application consistent backups of enterprise applications that are running on Azure or on-premises.

Microsoft Azure Backup Server requires an instance of Windows Server 2012 R2 that can run on Azure or on-premises. The preferred location depends on the scenario that you want to protect.



After the installation of Azure Backup Server, it needs to be registered with an Azure backup vault. The Azure Backup agent needs to be installed on the workload server that should be protected. Application-specific backups/restores can be configured afterward by using the MABS user interface. All backup data are encrypted with an encryption passphrase, similar to the approach in the previous section. Further details how to configure MABS are described here: <https://azure.microsoft.com/en-us/documentation/articles/backup-azure-microsoft-azure-backup/>.

7.3 Establishing secure remote access

Business Application operations require remote access for system administrators in order to fix operational issues or to extend, update, or reconfigure an application. Customers are facing the problem that desktops of the system administrators are not necessarily in the same virtual network as the servers that they have to administrate, and gaining remote access via public IP addresses is not appropriate for a lot of enterprises due to security reasons. Enterprises have different solutions in place such as dedicated administration networks that require an additional network interface card on every machine, jump servers with web-based RDP capabilities, or remote desktop services with Internet gateways and Multi-Factor Authentication that are difficult to set up, configure, and maintain.

Sometimes, especially in enterprise environments, firewalls prevent connecting via RDP to Azure Windows VMs over port 3389. Quite often, the only outgoing ports being open in the network are 80 and 443 for HTTP(S). Customers that have no explicit demand for securing remote access with Multi-Factor Authentication could consider web-based RDP clients that are offered by various third-party solutions. Customers with a demand for Multi-Factor Authentication should consider using Azure RemoteApp for gaining secure remote access.

Azure RemoteApp brings the functionality of the on-premises Microsoft RemoteApp program, backed by Remote Desktop Services, to Azure. Azure RemoteApp helps you provide secure, remote access to applications from many different user devices. Azure RemoteApp basically hosts nonpersistent Terminal Server sessions in the cloud, and you get to use them and share them with your users. With Azure RemoteApp you can share apps and resources with users on almost any device. Remote Apps uses corporate credentials from Azure AD, letting you ensure the security of apps and data, and you can combine it with Azure Multi-Factor Authentication.

The focus of this section is the usage of Azure Remote App to get remote access to a Windows or Linux virtual machine in a secure way. For this scenario it is best practice to use MFA for all Remote App users.

Azure RemoteApp lets you share apps and resources with users on any device. You do this by creating collections to hold the apps and resources, and then you share those collections with users. There are two different collection options, with different network and authentication options:

- **Cloud collections**


This type of collection is very quick to create. You can use one of the default images (Windows Server 2012 R2, Windows Server with Office Professional Plus 2013, or Windows Server with Office 365 Pro Plus) or you can use a custom image built from an Azure VM. You can also use your own Azure VNET to provide access into your on-

premises environment for data sharing or to use non-Windows authentication into resources. Authentication of users is done by using Azure Active Directory.


- **Hybrid collections**

These collections provide full access to on-premises network and Azure VNET. They include domain join access for apps and data. Remote applications can authentication against your on-premises Active Directory—they can then access resources in your domain.

After you create your RemoteApp collection, you need to publish the apps or resources that you want to make available for your users. The template images provided with your subscription only have a few apps published by default—to share the other apps, you need to publish them. For the remote access scenario, it is sufficient to publish the remote desktop connection for windows and a command line interface or other preferred tooling for Linux. You can publish a remote desktop connection in a generic manner, without using any command line parameter or you can publish it multiple times with command line parameters that are pointing directly to the servers that you want to grant access to. Due to the connectivity into the virtual network, it is not required to access any of the servers via a public IP address.



MyApplicationServer

 Veröffentlicht

%SYSTEMDRIVE%\Windows\system32\mstsc.exe

EIGENSCHAFTEN DES REMOTEAPP-PROGRAMMS

NAME

MyApplicationServer

BEFEHLSZEILENPARAMETER

/v:10.0.1.17

Before your users can see and use the apps in Azure RemoteApp, you have to grant them access to your collection. The different collection types support using different user identities for access to applications. For a hybrid collection of RemoteApp, you need to set up an Active Directory domain infrastructure on-premises and an Azure Active Directory tenant with Directory Integration and optionally single sign-on. In addition, you need to create some Active Directory objects in the on-premises directory.

For a cloud collection of RemoteApp, any user that has Azure Active Directory support identities can be granted user access to RemoteApp.

User accounts	Cloud	Hybrid
Microsoft Account	Yes	No

Azure Active Directory (Azure AD)		
Azure AD cloud only	Yes	No
AD Sync with password sync	Yes	Yes
AD Sync without password sync	Yes	No
AD Sync with AD FS	Yes	Yes
Third-party Azure-supported identity providers (for example, Ping)	Yes	Yes
Multi-Factor Authentication	Yes	Yes

One of the beauties of Azure RemoteApp is that you can access apps from any of your devices. The following operating systems are supported:

- Windows 7 Service Packs 1, 8, 8.1, 10
- Windows Phone 8.1
- iOS
- Mac OS X
- Android

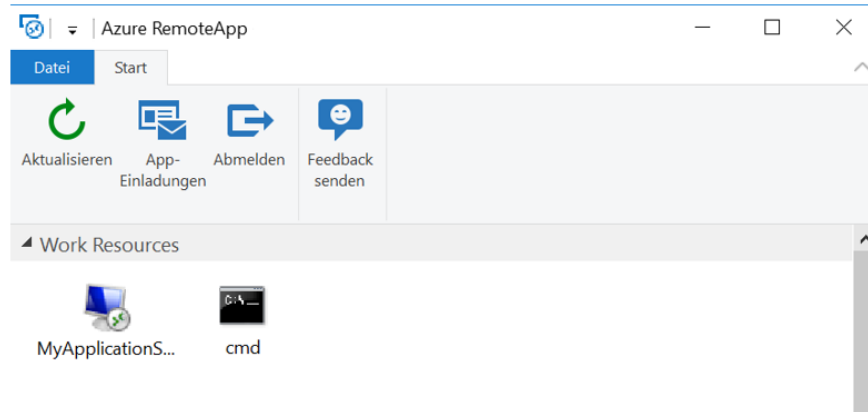
The following Windows Embedded thin clients are supported:

- Windows Embedded Standard 7
- Windows Embedded 8 Standard
- Windows Embedded 8.1 Industry Pro
- Windows 10 IoT Enterprise

The Azure Remote App Client can be downloaded here:

<https://www.remoteapp.windowsazure.com>.

After logging in to the remote app client you will see the apps that have been published for you.



7.4 Automating operational procedures

Microsoft Azure Automation provides a way for users to automate the manual, long-running, error-prone, and frequently repeated tasks that are commonly performed in a cloud and enterprise environment. It saves time and increases the reliability of regular administrative tasks and even schedules them to be automatically performed at regular intervals. You can automate processes using runbooks or automate configuration management using Desired State Configuration. A runbook is a set of tasks that performs some automated process in Azure Automation. It may be a simple process such as starting a virtual machine and creating a log entry, or you may have a complex runbook that combines other smaller runbooks to perform a complex process across multiple resources or even multiple clouds and on-premises environments.

Runbooks in Azure Automation are based on Windows PowerShell or Windows PowerShell Workflow, so they do anything that PowerShell can do. If an application or service has an API, then a runbook can work with it. If you have a PowerShell module for the application, then you can load that module into Azure Automation and include those cmdlets in your runbook. Azure Automation runbooks run in the Azure cloud and can access any cloud resources or external resources that can be accessed from the cloud. Using Hybrid Runbook Worker, runbooks can run in your local datacenter to manage local resources. The Runbook Gallery contains runbooks from Microsoft and the community that you can either use unchanged in your environment or customize for your own purposes. They are also useful as references to learn how to create your own runbooks.

Especially interesting in regard to this guide is the integration of Azure Automation into the Alerting of the Log Analytics Service. This allows an automated remediation of some issues that are typically coming up. For example, if you are running out of disk space you can add additional data disks or increase the size of the disks of your Azure VM in an automated fashion.

7.5 Managing IT services according to ITIL

The main focus of the operations layer is to carry out the business requirements that are defined at the service delivery layer (see section 9). Cloud service attributes for Azure IaaS cannot be achieved through technology alone; mature IT service management is still required.

The components of the operations layer include:

- **Change management**
Responsible for controlling the lifecycle of all changes. The primary objective is to implement beneficial changes with minimum disruption to the perception of continuous availability.
- **Service asset and configuration management**
Maintains information about the assets, components, and infrastructure needed to provide a service. Accurate configuration data for each component and its relationship to other components must be captured and maintained.
- **Release and deployment management**
Ensures that changes to a service are built, tested, and deployed with minimal disruption to the service or production environment. Change management provides the approval mechanism (determining what will be changed and why), but release and deployment management is the mechanism for determining how changes are implemented.
- **Knowledge management**
Involves gathering, analyzing, storing, and sharing information within an organization. Mature knowledge management processes are necessary to achieve a service provider's approach, and they are a key element of IT service management.
- **Incident and problem management**
Resolves disruptive, or potentially disruptive, events with maximum speed and minimum disruption. Problem management also identifies root causes of past incidents and seeks to identify and prevent, or minimize the impact of, future ones.
- **Request fulfillment**
Manages user requests for services. As the IT department adopts a service provider's approach, it should define available services in a service catalog based on business functionality.
- **Access management**
Denies access to unauthorized users while making sure that authorized users have access to needed services. Access management implements security policies that are defined by information security management at the service delivery layer.

- **Systems administration**

Performs the daily, weekly, monthly, and as-needed tasks that are required for system health. A mature approach to systems administration is required for achieving a service provider's approach and for driving predictability. The vast majority of systems administration tasks should be automated.

Azure doesn't provide a cloud-based Service Management tool that is comparable to the approach that was taken with the Operations Management Suite (OMS). System Center 2012 R2 Service Manager is the product in the System Center suite that covers the service management processes. The goal of System Center 2012 R2 Service Manager is to support IT service management in a broad sense. This includes implementing the Information Technology Infrastructure Library (ITIL) and Microsoft Operations Framework (MOF) processes such as change and incident management. Customers with another ITSM tooling will most likely continue with it and integrate it with OMS.

7.6 Recommendations for operating Azure IaaS Services

Recommendations for operating Azure IaaS Services	
Use nonpeak hours for backups	Schedule backups during nonpeak hours for VMs so that backup service gets IOPS for transferring data from customer storage account to backup vault.
Spread VMs into different backup schedules	Please make sure that in a policy VMs are spread from different storage accounts. We suggest that if the total number of disks stored in a single storage account from VMs is more than 20, spread the VMs into different backup schedules to get required IOPS during transfer phase of the backup.
Back up critical data to separate location	Ensuring all applications and mission-critical data is backed up at a separate secondary location other than the primary datacenter

8 Migrating existing services to Azure

Azure Site Recovery (ASR) is primarily a solution for disaster recovery from on-premises to Azure or from one on-premises location to another. ASR allows organizations to replicate their on-premises systems into Azure and keep them synchronized at all times. ASR is cross-compatible with all types of on-premises environments including Hyper-V, VMware, and physical servers, so no matter the source, ASR will be able to replicate it into Azure, and ASR will take care of any conversion that is required. Once the systems are replicated to Azure, you can schedule how often they synchronize with the on-premises environment so that the copy of the VM in Azure is always up to date.

Since the Azure Site Recovery replica is always up to date, there is always a complete, ready-to-go copy of the environment sitting in Azure. Instead of having to do a lengthy migration with long downtime windows, or having to rebuild servers from scratch, Azure allows you to perform a migration where you can use your ASR replica and perform a failover.

Just as you would want in a DR scenario, initiating a failover ASR means you can turn off the on-premises server and bring the cloud server online to replace it within a short period of time. The cloud replica will have all of the same settings and configurations, so with a quick DNS change, your system will now be running completely from the cloud. If you want to test the migration ahead of time, the test failover option allows you to perform the full failover process without disabling the production system. There is, of course, some configuration that is required ahead of time to ensure the networking, storage, and performance is all configured properly. However, all of this can be staged in advance for an easy transition.

As an added incentive to use ASR for migrations, Microsoft is providing Azure Site Recovery for free for the first 31 days. If you can complete the migration within that time, you will not pay anything for the replication software. You will only be charged for any compute and storage used during replication and after cutover.

8.1 Configuring virtual machine and application migrations

The configuration of virtual machine migrations is done in the Azure Portal by using the Recovery Services Vault. The detailed configuration is depending on the source infrastructure: Hyper-V site to Azure, Virtual Machine Manager to Azure, VMware to Azure. The recovery service vault needs to register the source site for the migration. Depending on the business continuity requirements, it is necessary to define settings for the copy frequency, recovery point retention, app consistent snapshot frequency, replication start time, and data encryption. It is not necessary to migrate all virtual machine from the source infrastructure. Virtual machines can

be selected. Furthermore, it is required to map source and target networks, choose the appropriate storage category for the migrated workload, and assign the compute power that is required for the migrated workload.

Once all settings have been configured it is recommended to perform a test migration before the planned migration will be executed.

It doesn't matter whether your virtual machines are Windows- or Linux-based, or running on VMware or Hyper-V VMs. Site Recovery integrates with Microsoft applications, including SharePoint, Exchange, Dynamics, SQL Server, and Active Directory. Microsoft also works closely with leading vendors including Oracle, SAP, IBM, and Red Hat to ensure their applications and services work well with ASR and Azure.

8.2 Mapping of networks and subnets

When customers are planning to migrate workloads, one of the key questions in their minds is how the virtual machine would be reachable after the failover is completed. ASR allows the administrator to choose the network to which a virtual machine would be connected to after failover. There are two choices for designing the network for the migrated VMs:

- **Use a different IP address range for the network at the Azure site.**

In this scenario the virtual machine will get a new IP address after failover, and the administrator would have to do a DNS update. This approach seems to be the most prevalent based on what we have seen. It takes the form of changing the IP address of every VM that is involved in the migration. A drawback of this approach requires the incoming network to "learn" that the application that was at IPx is now at IPy. Even if IPx and IPy are logical names, DNS entries typically have to be changed or flushed throughout the network, and cached entries in network tables have to be updated or flushed, therefore a downtime could be seen depending on how the DNS infrastructure has been set up.

- **Use same IP address range for the network at the Azure site.**

In a lot of scenarios administrators prefer to retain the IP addresses that they have on the primary site even after the migration. From a migration perspective, using fixed IP addresses appears to be the easiest method to implement, but it has a number of potential challenges that in practice make it the least popular approach. For the migration of virtual machines to Azure it is required to use subnet failover. In order to maintain the IP address space for the migration, it is possible to programmatically arrange for the router infrastructure to move the subnets from one site to another. In a migration scenario the subnets would move with the associated protected VMs. The

main drawback to this approach is that you have to move the whole subnet, which may be OK but it may affect the migration granularity considerations.

8.3 Planning and testing failover

Once the VMs are ready to be included in a failover strategy, Site Recovery can be used to monitor for failover events and then automate the process based on steps detailed in a failover plan. This can be as simple as bringing up a single server that has been migrated to Azure to orchestrating an entire environment transition that includes shutting down servers in the primary datacenter and bringing the migrated servers up in Azure. The startup sequence can even be specified.

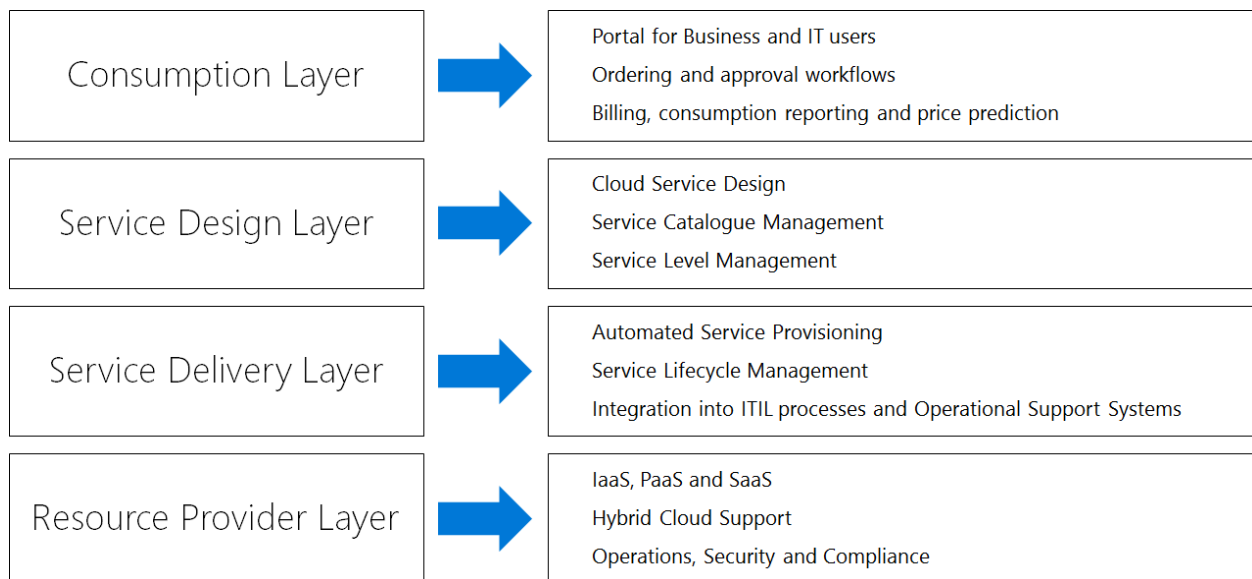
Of course a failover plan is only useful if it actually works, and the only way to ensure it works is to test it. Site Recovery makes this step possible by leveraging the flexibility of Azure Virtual Networks. The testing can even be put on a schedule. When a test is started, Site Recovery will spin up the failover VMs in an isolated network so they can be running at the same time as production. Users can then connect to the recovered servers and verify everything is working. When the test is complete, Site Recovery will tear down the VMs so there is no manual cleanup required. Depending on the workloads being tested, there may be some additional steps involved.

8.4 Recommendations for migrating existing services to Azure

Recommendations for migrating existing services to Azure	
Test before you migrate	Test applications in Azure before migration.
Automate the migration process	If you're failing over to Azure, for the best RTO we recommend that you automate all manual actions by integrating with Azure automation and recovery plans.
Check network bandwidth	Site Recovery supports a near-synchronous recovery point objective (RPO) when you replicate to Azure. Make sure that you have sufficient bandwidth between your datacenter and Azure.

9 Offering management for cloud-based services

All large IT organizations have some processes in place to enable their internal customers (for example, business departments) to obtain IT services and products that are available in the organization. These processes may vary based on the service needed, the department that provides it, or even the worker's location. A common vision is to have a unified service catalog for delivering services across the organization—some kind of a one-stop-shopping approach that enables customers to efficiently submit their requests. The typical logical layers for such an approach are as follows:



The consumption layer presents a service catalog to the consumer, with various configuration options (Application, Compute, RAM, Storage, Location, Network, SLA, etc.), including pricing information. Approval workflows are sometimes used to control costs and to govern the consumption of services.

The service layer is used for managing the service catalog, defining the SLAs (Availability, Performance, Time to React, etc.) that are associated with the provided services, and designing the service itself. A service consists not only of a virtual machine with an operating system. Enterprise IT departments are usually offering managed services for operating systems, databases, middleware products, web servers, or complex business applications that are based on multiple of these components. The services need to be connected to the monitoring, backup, and antivirus and managed according to ITIL standards.

The service provisioning as well as the integration with the operational support systems and the IT service management tooling is part of the service delivery layer. Those processes could be

automated, semi-automated, or executed manually with a defined service level for delivering the service to the consumer.

Finally, there is Resource Provider Layer that with IaaS-, PaaS-, and SaaS-based services is used as a foundation to create managed services. A large portion of enterprise IT is still focused on IaaS, but the obvious benefits of PaaS and SaaS are leading to a higher adoption of those services.

We observe that customers are sometimes aiming for a multivendor strategy on the resource provider layer called multicloud support or cloud brokerage. But there are a lot of caveats with such an approach:

- Each cloud provider has its own standards and technology. Adopting a multicloud approach will actually reduce the agility and the capabilities that customers get from the cloud services they are paying for. Higher level services like PaaS and SaaS differ heavily between the various providers, and on IaaS there is also a significant discrepancy in regards to network, security and resilience, operations, and lifecycle management. Using the lowest common denominator between the clouds has significant disadvantages for the business.
- The speed of innovation in Azure and the frequency of releasing new services is very challenging for any integration layer. Business departments want to benefit from the new capabilities, but the cloud brokering platform isn't supporting it in a timely manner.
- The complexity and lock that is introduced with such an integration layer is enormous.

Customers should rate for themselves the pros and cons of multicloud support. In the majority of situations, we feel that the selection of a trusted partner, with a cloud service offering that fits well into the enterprise architecture and with strong support for hybrid scenarios, might be the better approach.

9.1 Consuming services

The user self-service capability is an essential characteristic of cloud computing, and it must be present in any implementation. The intent is to permit users to approach a self-service capability and be presented with options available for provisioning. The capability may be basic (such as provisioning of a virtual machine with a predefined configuration), more advanced (such as allowing configuration options to the base configuration), or complex (such as implementing a platform capability or service).

The self-service capability is a critical business driver that allows members of an organization to become more agile in responding to business needs with IT capabilities that align and conform to internal business and IT requirements. The interface between IT and the business should be

abstracted to a well-defined, simple, and approved set of service options. The options should be presented as a menu in a portal or available from the command line. Businesses can select these services from the catalog, start the provisioning process, and be notified upon completion. They are charged only for the services they actually used.

The challenge is to find the right balance between the required configuration options to fulfil the business needs and the complexity to implement those options and provision the services accordingly. Adding all configuration options that Azure provides for a service to the service catalog doesn't make sense. It would end up in re-creating the Azure Portal, which isn't achievable for any organization. A common approach is to add the most important parameters that allow a basic automated provisioning of the service to the service catalog. Detailed configuration settings could be requested by the customer in a separate service request and could be performed by IT staff within a defined service level.

9.2 Provisioning of cloud services

The infrastructure for an application is typically made up of many components—maybe a virtual machine, storage and virtual network, or a web app, database, database server, and probably third-party services from the marketplace. These components are not considered as separate entities; instead you see them as related and interdependent parts of a single entity. The goal is to deploy, manage, and monitor them as a group. Azure Resource Manager enables customers to work with the resources in a solution as a group. You can deploy, update, or delete all of the resources of an application in a single, coordinated operation. Templates could be used for deployment and that template can work for different environments such as testing, staging, and production. Resource Manager also provides security, auditing, and tagging features to help you manage your resources after deployment.

A resource group is a container that holds related resources for an application. The resource group could include all of the resources for an application, or only those resources that are logically grouped together. The service designer decides how to allocate resources to resource groups based on what makes the most sense for the organization.

With Resource Manager, application designers can create a simple template (in JSON format) that defines deployment and configuration of entire application. This template is known as a Resource Manager template and provides a declarative way to define deployment. By using a template, you can repeatedly deploy the application throughout the app lifecycle and have confidence that resources are deployed in a consistent state.

Within the template, you define the infrastructure for your app, how to configure that infrastructure, and how to publish your app code to that infrastructure. You do not need to worry about the order for deployment because Azure Resource Manager analyzes dependencies

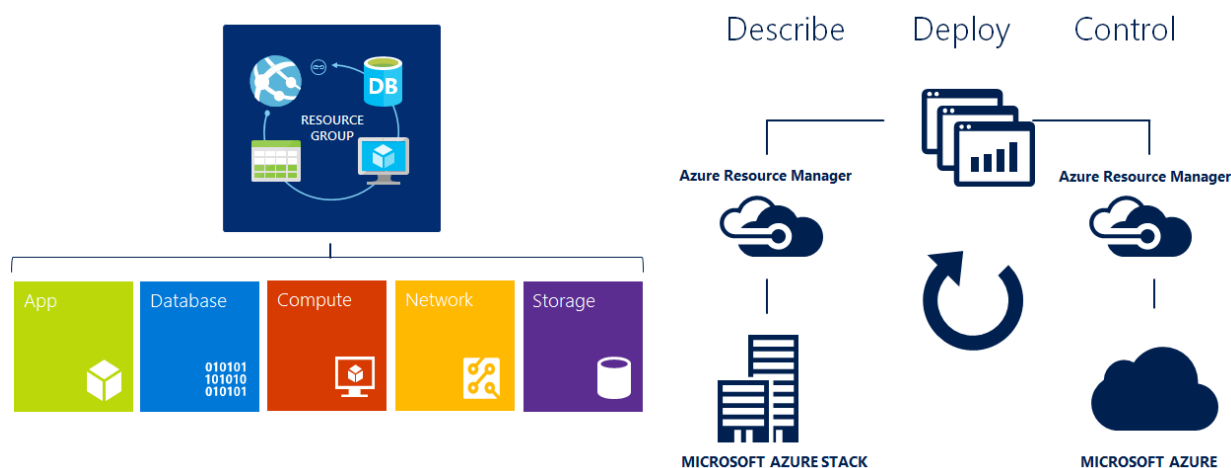
to ensure resources are created in the correct order. There is no need to define your entire infrastructure in a single template. Often, it makes sense to divide the deployment requirements into a set of targeted, purpose-specific templates that are linked together.

Templates allow the specification of parameters for customization and flexibility in deployments. Those parameters should fit to the configuration options in the service catalog. The fact that resource manager orchestrates the deployment of multiple components supports a definition of application-specific instead of infrastructure-specific parameters. Example: For a deployment of a Big Data Service you may ask the customer for parameters such as the amount of master and data nodes, the required performance, and the amount of data that should be handled. All the logic for provisioning multiple virtual machines, load balancers, network settings, application installation, and configuration, etc., could be baked into the template.

There are a huge number of templates available on GitHub, <https://github.com/Azure/azure-quickstart-templates>, which provides a real productivity boost for any IT organization.

Many enterprises might choose to keep some applications on-premises. Perhaps they are based on nonstandard systems or legal constraints don't allow to move it to the cloud. This results in the situation where most enterprise customers are running a hybrid cloud scenario.

Microsoft Azure Stack is designed for this approach and allows customers to run the majority of Azure services within their own datacenter. Azure Stack brings with it a huge innovation to the hybrid-cloud ecosystem. Before the evolution of Azure Stack, the Azure cloud resource and on-premises resource management were different, requiring separate deployment scripts and potentially significant code changes to move systems either from Azure to on-premises or on-premises to Azure. With the release of Microsoft Azure Stack, it is now possible to use the same deployment and development tools to build and deploy applications and infrastructure that reside either in the Azure Cloud or within the enterprise on-premises datacenter.



9.3 Metering consumption per application

Customers require the ability to accurately predict and manage their application costs. As they move from a Capex to an Opex model, they also need the ability to do Showback versus Chargeback analysis, as well as provide mode fidelity in estimation and billing, especially for large deployments. The Azure Resource Usage and Rate Card APIs address these needs, by enabling new insights into the consumption of Azure resources.

The Azure Usage API allows subscribers to programmatically pull in usage data to gain insights into their consumption. The granularity (hourly usage information) and resource metadata information available through the API provides the necessary dataset to support flexible Showback or Chargeback models. Supported features are:

- **Azure Role-Based Access Control**
Customers can configure their access policies on the Azure Portal or through Azure PowerShell cmdlets to specify which users or applications can get access to the subscription's usage data.
- **Hourly or daily aggregations**
Callers can specify whether they want their Azure usage data in hourly buckets or daily buckets. The default is daily.
- **Instance metadata provided**
Instance-level details such as the fully qualified resource uri along with the resource group information and resource tags will be provided in the response. This will help customers deterministically and programmatically allocate usage by the tags, for use cases like cross-charging.
- **Resource metadata provided**
Resource details such as the meter name, meter category, meter subcategory, unit, and region will also be passed in the response, to give the callers a better understanding of what was consumed.
- **Usage for all offer types**
Usage data will be accessible for all offer types, including Pay-as-you-go, MSDN, Monetary commitment, Monetary credit, and Enterprise Agreement, among others.

The data available through the Azure Usage API includes not only consumption information but also resource metadata including any tags associated with it. Tags provide an easy way to organize application resources, but in order to be effective you need to ensure that:

- Tags are correctly applied to the application resources at provisioning time.

- Tags are properly used on the Showback/Chargeback process to tie the usage to the organization's account structure.

9.4 Billing and price prediction

A proper price prediction for a potentially consumed service from the service catalog is a common demand in most organizations. Managers require this information to approve large cloud application deployments. Price prediction requires detailed knowledge about the Azure elements that a service consists of (defined at design time) along with estimated pricing information for each.

The Azure Resource RateCard API delivers a list of available Azure resources and pricing information. Features include:

- **Azure Role-Based Access Control**
Customers can configure their access policies on the Azure Portal or through Azure PowerShell cmdlets to specify which users or applications can get access to the RateCard data.
- **Support for Pay-as-you-go, MSDN, Monetary commitment, and Monetary credit offers (EA not supported)**
This API provides Azure offer-level rate information versus subscription-level. The caller of this API must pass in the offer information to get resource details and rates. As EA offers have customized rates per enrollment, we are unable to provide the EA rates at this time.

Enterprise Agreement customers can use another API that allows usage access price sheet and other billing information in format of CSV and JSON from API. Enterprise Administrators control access to the API through access keys. Available reports through the API are:

- **Enrollment Summary CSV/JSON File**
This report contains information regarding the enrollment summary for the month. The report will have the same data content as the Balance and Charge report available on the EA portal under the Download Usage Data section.
- **Usage and Billing Details CSV/JSON File**
This report will have detailed information regarding service usage and billing. The report will have the same data content as the Usage Detail report available in the EA portal under the Download Usage Data section.
- **Marketplace StoreCharge CSV/JSON File**
This report contains the same data content as the StoreCharge downloaded from EA portal under the Download Usage Data section.

- **Price Sheet CSV/JSON File**

The report contains the same data content as the Price sheet downloaded from the EA portal under the Download Usage Data section.

There are also partner solutions available that make use of those APIs and provide an integrated experience:

- [Microsoft Azure Usage and RateCard APIs Enable Cloudyn to Provide ITFM for Customers](#) describes the integration experience offered by Azure Billing API partner Cloudyn. This article provides detailed coverage of its experiences, including a short video that shows how Azure customers can use Cloudyn and the Azure Billing APIs to gain insights from their Azure consumption data.
- [Cloud Cruiser and Microsoft Azure Billing API Integration](#) describes how Cloud Cruiser's Express for Azure Pack works directly from the WAP portal, enabling customers to seamlessly manage both the operational and financial aspects of their Microsoft Azure private or hosted public cloud from a single user interface.

9.5 Managing the lifecycle

After a service has been provisioned it needs to be managed along the entire lifecycle. This includes configuration changes, applying updates, performing actions, and decommissioning the service.

Maintenance of solutions in Azure is largely dependent on the services that are consumed within Azure (PaaS or IaaS). PaaS Services are fully managed by Microsoft, but Azure IaaS virtual machines have the requirement to be maintained by the customer. Microsoft does not provide a centralized patch management offering for the guest operating system of IaaS virtual machine outside of currently shipping patch management solutions such as Windows Server Update Services (WSUS) and System Center Configuration Manager. The underlying fabric hardware, virtualization, and service layers are managed by Azure. Keeping up-to-date with Microsoft updates for Windows-based virtual machines is critical to ensure that a proper security posture is maintained for these systems. Microsoft updates should be applied to Azure IaaS virtual machines in a similar way that updates are applied to other customer environments.

When updating from on-premises or public Microsoft update servers, the updates source location is largely driven by the Azure network design decisions and customer configurations—like any other Windows-based virtual machine.

For organizations with small environments, or organizations that have not invested in a patching infrastructure (such as System Center Configuration Manager or a similar third-party tool), WSUS can provide a basic-level patch management infrastructure. However, the virtual machines need

to be configured to utilize the WSUS instance like any other Windows-based system in the enterprise.

Like Windows Update, Windows Server Update Services (WSUS) can be utilized to patch Azure IaaS virtual machines for customers who want to have a higher degree of control over patch distribution, release, and reporting. If an organization has an existing WSUS topology, a recommended approach is to deploy an additional WSUS server within the organization's Azure subscription and joined to the WSUS hierarchy. Optionally, this additional server can be configured to be a content store, such that Azure virtual machines download content from this new server.

System Center Configuration 2012 R2 Manager can provide services including installing applications and updating management, and other system configuration tasks. This is particularly attractive in hybrid scenarios where customers may have significant existing investments in Configuration Manager packaging, software update groups, and so on. Microsoft Azure presents additional configurations that should be considered, such as updating location settings, boundaries, and client authentication.

If Configuration Manager is going to be used, we recommend that organizations leverage an existing Configuration Manager infrastructure, if available. The Configuration Manager architecture (such as primary sites and distribution points) can be an involved and generally a separate engagement beyond an Azure scope of work.

Besides the above-mentioned maintenance tasks, we see demands from the business to support typical lifecycle actions in the service catalog such as:

- Start the service
- Stop the service
- Restart the service
- Create an ad-hoc backup of the service
- Decommission the service

Organizations have to make up their minds whether those actions might be executed in an automated fashion by using the Azure APIs or be performed manually within a defined service level.

9.6 Recommendations for cloud service provisioning

Recommendations for cloud service provisioning

Use Resource Manager	Azure has two different deployment models: Resource Manager and Classic. Microsoft recommends using Resource Manager for new deployments.
Follow reference architectures	Running a Single Windows VM on Azure Running multiple Windows VM instances on Azure (single tier, Internet-facing) Running Windows VMs for an N-tier architecture on Azure Adding reliability to an N-tier architecture on Azure (Windows VMs) Running Windows VMs in multiple datacenters on Azure Extending an on-premises network to Azure using a site-to-site virtual private network Implementing a highly available hybrid network architecture in Azure by using failover between ExpressRoute and VPN gateway Implementing a hybrid network architecture with Azure ExpressRoute

About the Author



Joachim Hafner is a Cloud Solution Architect at Microsoft Germany helping clients to integrate the Azure platform into their enterprise architecture and provides guidance for architecting modern cloud applications based on Azure technologies. Before he joined Microsoft he was working as a Senior Enterprise Architect for one of the largest IT service providers with a strong focus on hybrid cloud strategies.