

Domain 1: Introduction to Cloud

Infrastructure: Describe cloud concepts

Describe cloud computing

1. Introduction to Microsoft Azure Fundamentals

<https://learn.microsoft.com/en-us/training/modules/describe-cloud-compute/1-introduction-microsoft-azure-fundamentals>


Introduction to Microsoft Azure Fundamentals

Completed

Microsoft Azure is a cloud computing platform with an ever-expanding set of services to help you build solutions that meet your technical goals. Azure services support everything from simple to complex. You can host simple web services for internet-facing apps, run fully virtualized computers for custom software solutions, or use cloud-based services like remote storage, database hosting, and centralized account management. Azure also offers capabilities in artificial intelligence (AI) and Internet of Things (IoT).

In this series, you cover cloud computing basics, get introduced to some of the core services provided by Microsoft Azure, and learn more about the governance and compliance services that

you can use.



What is Azure Fundamentals?

Azure Fundamentals is a series of three learning paths that familiarize you with Azure and its many services and features and one learning path that gives you a chance to explore Azure and build solutions to realistic IT challenges.

Whether you're interested in compute, networking, or storage services; learning about cloud security best practices; or exploring governance and management options, think of Azure Fundamentals as your curated guide to Azure.

Azure Fundamentals includes guided content and knowledge checks that reinforce key concepts.

Technical IT experience isn't required; however, having general IT knowledge will help you get the most from your learning experience.

Why should I take Azure Fundamentals?

If you're just beginning to work with the cloud, or if you already have cloud experience, Azure Fundamentals provides you with everything you need to get started.

No matter your goals, Azure Fundamentals has something for you. You should take this course if you:

- Have general interest in Azure or in cloud computing
- Want to earn official certification from Microsoft (AZ-900)

The Azure Fundamentals learning path series can help you prepare for Exam AZ-900: Microsoft Azure Fundamentals. This exam includes three knowledge domain areas:

AZ-900 Domain Area	Weight
Describe cloud concepts	25-30%
Describe Azure architecture and services	35-40%
Describe Azure management and governance	30-35%

Each domain area maps to a learning path in Azure Fundamentals. The percentages shown indicate the relative weight of each area on the exam. The higher the percentage, the more questions that part of the exam will contain. Be sure to read the exam page for specifics about what skills are covered in each area.

This training helps you develop a broad understanding of Azure.

2. Introduction to cloud computing

<https://learn.microsoft.com/en-us/training/modules/describe-cloud-compute/2-introduction-cloud-compute>

Introduction to cloud computing

Completed

In this module, you're introduced to general cloud concepts. You start with an introduction to the cloud in general. Then you dive into concepts like shared responsibility, different cloud models, and explore the unique pricing method for the cloud.

If you're already familiar with cloud computing, this module may be largely review for you.

Learning objectives

After completing this module, you'll be able to:

- Define cloud computing.
- Describe the shared responsibility model.
- Define cloud models, including public, private, and hybrid.
- Identify appropriate use cases for each cloud model.
- Describe the consumption-based model.
- Compare cloud pricing models.

3. What is cloud computing

What is cloud computing

Completed

Cloud computing is the delivery of computing services over the internet. Computing services include common IT infrastructure such as virtual machines, storage, databases, and networking. Cloud services also expand the traditional IT offerings to include things like Internet of Things (IoT), machine learning (ML), and artificial intelligence (AI).

Because cloud computing uses the internet to deliver these services, it doesn't have to be constrained by physical infrastructure the same way that a traditional datacenter is. That means if you need to increase your IT infrastructure rapidly, you don't have to wait to build a new datacenter—you can use the cloud to rapidly expand your IT footprint.


Practical example

Suppose a retail team expects high traffic during a seasonal launch. Instead of buying and configuring extra physical servers months in advance, they can deploy additional compute capacity in the cloud for the launch period and scale down afterward. This approach improves agility and helps align spending to actual demand.

At a fundamentals level, cloud computing changes infrastructure planning from long procurement cycles to on-demand provisioning. Teams can test faster, recover faster, and adapt capacity as requirements change.

Cloud platforms also offer global reach, so teams can place services closer to users and design for regional resilience without building multiple physical datacenters.

This short video provides a quick introduction to cloud computing.



4. Describe the shared responsibility model

Describe the shared responsibility model

Completed

You may have heard of the shared responsibility model, but you may not understand what it means or how it impacts cloud computing.

How responsibilities shift in cloud

Start with a traditional on-premises datacenter. Your team is responsible for maintaining the physical space, ensuring security, and maintaining or replacing the servers if anything happens. The IT department is responsible for maintaining all the infrastructure and software needed to keep the datacenter up and running. They're also likely to be responsible for keeping all systems patched and on the correct version.

With the shared responsibility model, these responsibilities get shared between the cloud provider and the consumer. Physical security, power, cooling, and network connectivity are the responsibility of the cloud provider. The consumer isn't collocated with the datacenter, so it wouldn't make sense for the consumer to have any of those responsibilities.

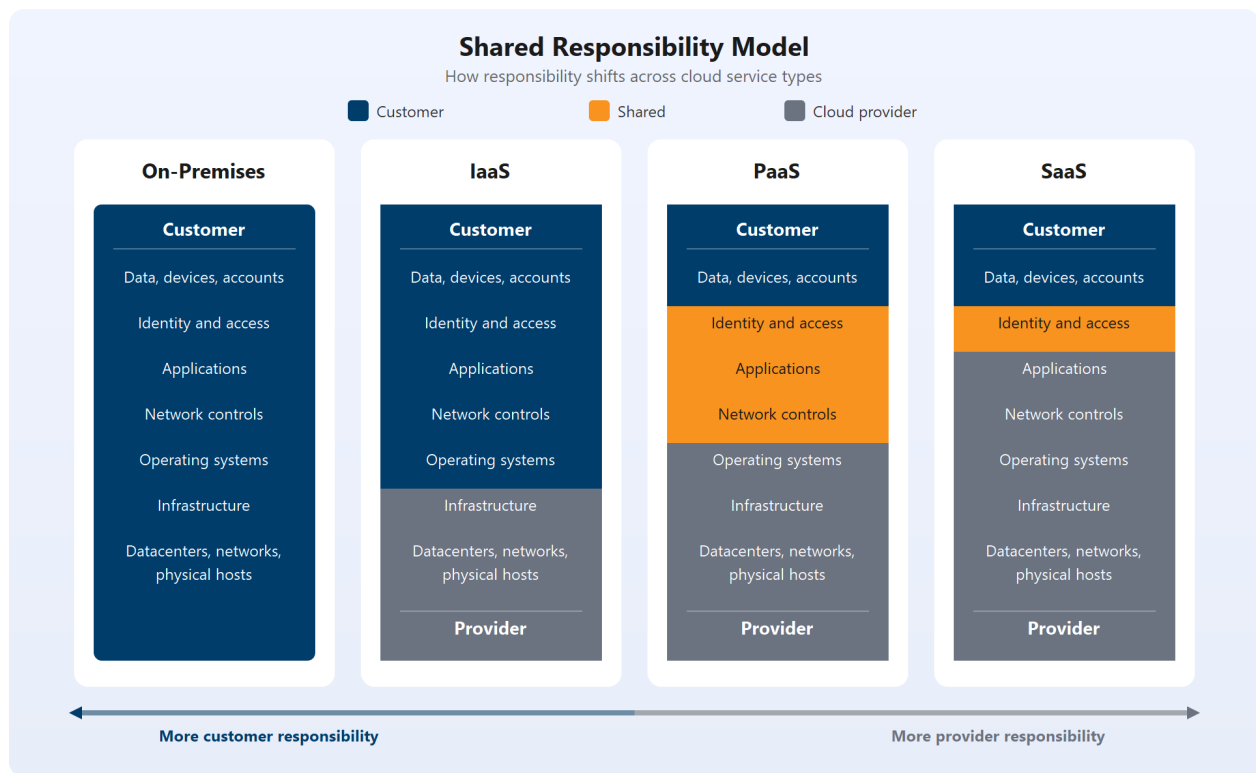
At the same time, the consumer is responsible for the data and information stored in the cloud. (You wouldn't want the cloud provider to be able to read your information.) The consumer is also responsible for access security, meaning you only give access to those who need it.

Then, for some things, the responsibility depends on the situation. If you're using a cloud SQL database, the cloud provider would be responsible for maintaining the actual database. However, you're still responsible for the data that gets ingested into the database. If you deployed a virtual machine and installed an SQL database on it, you'd be responsible for database patches and updates, as well as maintaining the data and information stored in the database.

With an on-premises datacenter, you're responsible for everything. With cloud computing, those responsibilities shift. The shared responsibility model is heavily tied into the cloud service types (covered later in this learning path): infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS). IaaS places the most responsibility on the consumer, with the cloud provider being responsible for the basics of physical security, power, and connectivity. On the other end of the spectrum, SaaS places most of the responsibility with the cloud provider. PaaS, being a middle ground between IaaS and SaaS, rests somewhere in the middle and evenly distributes responsibility between the cloud provider and the consumer.

Responsibility by service model

The following diagram highlights how the Shared Responsibility Model informs who is responsible for what, depending on the cloud service type.



When using a cloud provider, you'll always be responsible for:

What always stays with you

- The information and data stored in the cloud
- Devices that are allowed to connect to your cloud (cell phones, computers, and so on)
- The accounts and identities of the people, services, and devices within your environment

The cloud provider is always responsible for:

What the provider always owns

- The physical datacenter
- The physical network
- The physical hosts

Your service model will determine responsibility for things like:

What depends on the service type

- Operating systems
- Network controls

- Applications
- Identity and access
- Infrastructure

For example, identity and access is shared in PaaS and SaaS—you manage your own users, roles, and policies, while the provider runs the authentication platform (such as Microsoft Entra ID).

Infrastructure, on the other hand, shifts entirely to the provider as soon as you move off-premises to IaaS.


5. Define cloud models

<https://learn.microsoft.com/en-us/training/modules/describe-cloud-compute/5-define-cloud-models>

Define cloud models

Completed

What are cloud models? The cloud models define the deployment type of cloud resources. The three main cloud models are: private, public, and hybrid.



Private cloud

A private cloud is a cloud environment used by a single entity. It evolved naturally from the traditional datacenter model, delivering IT services over the internet while keeping resources dedicated to one organization. Private cloud provides much greater control for your IT team. However, it also comes with greater cost and fewer of the benefits of a public cloud deployment. A private cloud may be hosted from your on-site datacenter, or in a dedicated datacenter offsite, potentially even by a third party.

Public cloud

A public cloud is built, controlled, and maintained by a third-party cloud provider. With a public cloud, anyone that wants to purchase cloud services can access and use resources. The general public availability is a key difference between public and private clouds.

Hybrid cloud

A hybrid cloud is a computing environment that uses both public and private clouds in an inter-connected environment. A hybrid cloud environment can be used to allow a private cloud to surge for increased, temporary demand by deploying public cloud resources. Hybrid cloud can be used to provide an extra layer of security. For example, users can flexibly choose which services to keep in public cloud and which to deploy to their private cloud infrastructure.

The following table highlights a few key comparative aspects between the cloud models.

Public cloud	Private cloud	Hybrid cloud
No capital expenditures to scale up	You have complete control over resources and security	Provides the most flexibility
Applications can be quickly provisioned and deprovisioned	Data isn't collocated with other tenants' data	You determine where to run your applications
You pay only for what you use	Hardware must be purchased for startup and maintenance	You control security, compliance, or legal requirements
You don't have complete control over resources and security	You are responsible for hardware maintenance and updates	

Multicloud

A fourth, and increasingly likely scenario is a multicloud scenario. In a multicloud scenario, you use multiple public cloud providers. Maybe you use different features from different cloud providers. Or maybe you started your cloud journey with one provider and are in the process of migrating to a different provider. Regardless, in a multicloud environment you deal with two (or more) public cloud providers and manage resources and security in both environments.

Azure Arc

Azure Arc is a set of technologies that helps manage your cloud environment. Azure Arc can help manage your cloud environment whether it's a public cloud solely on Azure, a private cloud in your datacenter, a hybrid configuration, or even a multicloud environment running on multiple cloud providers at once.

Azure VMware Solution

What if you're already established with VMware in a private cloud environment but want to migrate to a public or hybrid cloud? Azure VMware Solution lets you run your VMware workloads in Azure with seamless integration and scalability.

6. Describe the consumption-based model

<https://learn.microsoft.com/en-us/training/modules/describe-cloud-compute/6-describe-consumption-based-model>

Describe the consumption-based model

Completed

Cloud computing operates on a consumption-based model. You pay for the IT resources you use, and nothing more. Instead of buying and maintaining your own datacenter infrastructure, you rent compute power and storage and release those resources when you're done.

In traditional IT budgeting, you may hear the terms capital expenditure (CapEx) and operational expenditure (OpEx). CapEx refers to up-front spending on physical infrastructure like servers, network hardware, and datacenter space. OpEx refers to ongoing spending on services over time. Because you pay for cloud services as you consume them, cloud computing is classified as an operational expense.

The consumption-based model offers several key benefits:

- No upfront costs for hardware or datacenter infrastructure.
- No need to purchase and manage capacity that may go underutilized.
- The ability to add resources when demand increases.
- The ability to release resources when demand decreases.

Capacity planning: traditional vs. cloud

Capacity Planning



With a traditional datacenter, you estimate future resource needs in advance. Overestimate, and you overspend on infrastructure that sits idle. Underestimate, and your applications suffer degraded performance. Fixing the problem means ordering, installing, and provisioning additional hardware, power, and cooling.

In a cloud-based model, you adjust resources to match actual demand. Add virtual machines when you need more capacity; remove them when demand drops. You pay only for what you use, not for idle capacity. In practice, you can scale out during peak demand and scale back in when traffic subsides.

Cloud pricing

Cloud providers use a pay-as-you-go pricing model. You typically pay only for the services you consume, which helps you:

- Plan and manage operating costs.
- Run infrastructure more efficiently.
- Scale as workload needs change.

The cloud provider maintains the underlying infrastructure, including power, cooling, hardware, and networking, so you can focus on solving business problems and delivering new capabilities to your users.

7. Module assessment

Module assessment

Completed

8. Summary

<https://learn.microsoft.com/en-us/training/modules/describe-cloud-compute/8-summary>

Summary

Completed

In this module, you learned about general cloud concepts. You started with an understanding of what cloud computing is. You also learned about the shared responsibility model and how you and your cloud provider share the responsibility of keeping your information in the cloud secure. You covered the differences between the cloud models (public, private, hybrid, and multicloud). Then, you wrapped up with a unit on how the cloud shifts IT spend from a capital expense to an operational expense.

Learning objectives

You should now be able to:

- Define cloud computing.
- Describe the shared responsibility model.
- Define cloud models, including public, private, and hybrid.
- Identify appropriate use cases for each cloud model.
- Describe the consumption-based model.
- Compare cloud pricing models.

Additional resources

The following resources provide more information on topics in this module or related to this module.

- [Shared responsibility model](#) - The sharing of responsibilities for the cloud between you and your cloud provider.

- [Introduction to Azure VMware Solution](#) - A Microsoft Learn course that dives deeper into Azure VMware Solution.
- [Introduction to Azure hybrid cloud services](#) - A Microsoft Learn course that explains hybrid cloud in greater detail.

Explore with Copilot

Tip

Try one of these prompts in Copilot Chat:

- "Explain shared responsibility across on-premises, IaaS, PaaS, and SaaS using one concrete business application example."
- "Compare public, private, hybrid, and multicloud models, then recommend the best fit for three different organizational scenarios."
- "Create a financial comparison of CapEx versus OpEx for moving a workload to a consumption-based cloud model."

Describe the benefits of using cloud services

1. Introduction

<https://learn.microsoft.com/en-us/training/modules/describe-benefits-use-cloud-services/1-introduction>

Introduction

Completed

In this module, you're introduced to some of the benefits that cloud computing offers. You learn how cloud computing can help you meet variable demand while providing a good experience for your customer. You also learn about security, governance, sustainability, and overall manageability in the cloud.

Learning objectives

After completing this module, you'll be able to:

- Describe the benefits of high availability and scalability in the cloud.
- Describe the benefits of reliability and predictability in the cloud.
- Describe the benefits of security and governance in the cloud.
- Describe the benefits of manageability in the cloud.
- Describe sustainability considerations and benefits in the cloud.

2. Describe the benefits of high availability and scalability in the cloud

Describe the benefits of high availability and scalability in the cloud


Completed

When building or deploying a cloud application, two of the biggest considerations are uptime (or availability) and the ability to handle demand (or scale).

High availability

High Availability

Maximum uptime, regardless of disruptions




When you're deploying an application, a service, or any IT resources, it's important the resources are available when needed. High availability focuses on ensuring maximum availability, regardless of disruptions or events that may occur.

When you're architecting your solution, you need to account for service availability guarantees. Azure is a highly available cloud environment with uptime guarantees depending on the service. These guarantees are part of the service-level agreements (SLAs).

This short video describes Azure SLAs in more detail.

Scalability



Another major benefit of cloud computing is the scalability of cloud resources. Scalability refers to the ability to adjust resources to meet demand. If you suddenly experience peak traffic and your systems are overwhelmed, the ability to scale means you can add more resources to better handle the increased demand.

The other benefit of scalability is that you aren't overpaying for services. Because the cloud is a consumption-based model, you only pay for what you use. If demand drops off, you can reduce your resources and thereby reduce your costs.

Scaling generally comes in two varieties: vertical and horizontal. Vertical scaling is focused on increasing or decreasing the capabilities of resources. Horizontal scaling is adding or subtracting the number of resources.

Vertical scaling

With vertical scaling, if you were developing an app and you needed more processing power, you could vertically scale up to add more CPUs or RAM to the virtual machine. Conversely, if you realized you had over-specified the needs, you could vertically scale down by lowering the CPU or RAM specifications.

Horizontal scaling

With horizontal scaling, if you suddenly experienced a steep jump in demand, your deployed resources could be scaled out (either automatically or manually). For example, you could add additional virtual machines or containers. In the same manner, if there was a significant drop in demand, deployed resources could be scaled in (either automatically or manually).

3. Describe the benefits of reliability and predictability in the cloud

<https://learn.microsoft.com/en-us/training/modules/describe-benefits-use-cloud-services/3-reliability-predictability-cloud>

Describe the benefits of reliability and predictability in the cloud


Completed

Reliability and predictability are two crucial cloud benefits that help you develop solutions with confidence.

Reliability

Reliability

Recover from failures and continue to function



Decentralized design across regions

If one region fails, others keep running

Pillar of the Azure Well-Architected Framework


Reliability is the ability of a system to recover from failures and continue to function. It's also one of the pillars of the Microsoft Azure Well-Architected Framework.

The cloud, by virtue of its decentralized design, naturally supports a reliable and resilient infrastructure. With a decentralized design, the cloud enables you to have resources deployed in regions around the world. With this global scale, even if one region has a catastrophic event other regions are still up and running. You can design your applications to automatically take advantage of this increased reliability. In some cases, your cloud environment itself will automatically shift to a different region for you, with no action needed on your part. You'll learn more about how Azure leverages global scale to provide reliability later in this series.

Predictability

Predictability

Move forward with confidence



Predictability in the cloud lets you move forward with confidence. Predictability can be focused on performance predictability or cost predictability. Both performance and cost predictability are heavily influenced by the Microsoft Azure Well-Architected Framework. Deploy a solution built around this framework and you have a solution whose cost and performance are predictable.

Performance

Performance predictability focuses on predicting the resources needed to deliver a positive experience for your customers. Autoscaling, load balancing, and high availability are just some of the cloud concepts that support performance predictability. If you suddenly need more resources, autoscaling can deploy additional resources to meet the demand, and then scale back when the demand drops. Or if the traffic is heavily focused on one area, load balancing will help redirect some of the overload to less stressed areas.

Cost


Cost predictability is focused on predicting or forecasting the cost of the cloud spend. With the cloud, you can track your resource use in real time, monitor resources to ensure that you're using them in the most efficient way, and apply data analytics to find patterns and trends that help better plan resource deployments. By operating in the cloud and using cloud analytics and information, you can predict future costs and adjust your resources as needed. You can even use tools like the Azure Pricing Calculator to get an estimate of potential cloud spend.

4. Describe the benefits of security and governance in the cloud

Describe the benefits of security and governance in the cloud

Completed

Whether you're deploying infrastructure as a service or software as a service, cloud features support governance and compliance.



Tools such as templates help ensure that deployed resources meet your technical standards and regulatory requirements. As standards change, you can update resources at scale. Cloud-based auditing helps flag resources that are out of compliance with your baseline and provides mitigation strategies. Depending on your operating model, software patches and updates may also be applied automatically, which helps with both governance and security.

On the security side, you can find a cloud solution that matches your security needs. If you want maximum control of security, infrastructure as a service provides you with physical resources but lets you manage the operating systems and installed software, including patches and maintenance. If you want patches and maintenance taken care of automatically, platform as a service or software as a service deployments may be the best cloud strategies for you.

Cloud providers are typically well suited to handle things like distributed denial of service (DDoS) attacks, making your network more robust and secure.

By establishing a good governance footprint early, you can keep your cloud footprint updated, secure, and well managed.

5. Describe the benefits of manageability in the cloud


<https://learn.microsoft.com/en-us/training/modules/describe-benefits-use-cloud-services/5-manageability-cloud>

Describe the benefits of manageability in the cloud

Completed

A major benefit of cloud computing is the manageability options. There are two types of manageability for cloud computing, and both are excellent benefits.

Management of the cloud




Management of the cloud speaks to managing your cloud resources. In the cloud, you can:

- Automatically scale resource deployment based on need.
- Deploy resources based on a preconfigured template, removing the need for manual configuration.
- Monitor the health of resources and automatically replace failing resources.

- Receive automatic alerts based on configured metrics, so you're aware of performance in real time.

Management in the cloud



Management in the cloud speaks to how you're able to manage your cloud environment and resources. You can manage these:

- Through a web portal.
- Using a command line interface.
- Using APIs.
- Using PowerShell.

For example, an operations team can deploy resources from templates, monitor health in the portal, and automate recurring tasks with CLI or PowerShell scripts. This combination reduces manual effort and helps maintain consistent configurations.


6. Describe sustainability considerations in the cloud

Describe sustainability considerations in the cloud

Completed

Cloud computing can support sustainability goals when teams actively optimize how resources are deployed and used.

Sustainability in the Cloud



Right-size, automate, optimize, and monitor — sustainability in action

Why cloud can improve efficiency

Cloud providers operate at large scale, which can improve resource utilization compared to many isolated on-premises environments. In Azure, you can also reduce waste by matching deployed resources to actual demand.

Sustainability-aligned cloud practices

Examples of practices that support sustainability and cost efficiency include:

- Scaling resources down when demand decreases
- Turning off or deallocating resources that are not in use
- Choosing efficient services and configurations to reduce overprovisioning
- Using governance and monitoring to track usage trends and optimize deployments over time

For example, a development environment that runs only during business hours can be automatically shut down overnight and on weekends. This practice reduces unnecessary consumption while still meeting team needs.

At a fundamentals level, sustainability in the cloud is closely tied to good operational habits: right-size, automate, monitor, and continuously optimize.

7. Module assessment

<https://learn.microsoft.com/en-us/training/modules/describe-benefits-use-cloud-services/6-knowledge-check>

Module assessment

Completed

8. Summary

<https://learn.microsoft.com/en-us/training/modules/describe-benefits-use-cloud-services/7-summary>

Summary

Completed

In this module, you learned about some of the benefits of operating in the cloud. You learned about high availability and reliability, and how those work to keep your applications running. You also learned about how the cloud can provide a more secure environment and support sustainability goals. Finally, you learned that the cloud provides a highly manageable environment for your resources.

Learning objectives

You should now be able to:

- Describe the benefits of high availability and scalability in the cloud.
- Describe the benefits of reliability and predictability in the cloud.
- Describe the benefits of security and governance in the cloud.
- Describe the benefits of manageability in the cloud.
- Describe sustainability considerations and benefits in the cloud.

Additional resources

The following resources provide more information on topics in this module or related to this module.

- [Build great solutions with the Microsoft Azure Well-Architected Framework](#) - A Microsoft Learn course that introduces you to the Microsoft Azure Well-Architected Framework.

Explore with Copilot

Tip

Try one of these prompts in Copilot Chat:

- "Evaluate an on-premises system and map cloud improvements for availability, scalability, reliability, security, and manageability."
- "Design a KPI dashboard to measure whether cloud adoption is delivering the expected technical and business benefits."
- "Analyze tradeoffs where cloud benefits may be limited, including cost, governance complexity, and sustainability constraints."

Describe cloud service types

1. Introduction

<https://learn.microsoft.com/en-us/training/modules/describe-cloud-service-types/1-introduction>

Introduction

Completed

In this module, you're introduced to cloud service types. You learn how each cloud service type determines the flexibility you have with managing and configuring resources. You understand how the shared responsibility model applies to each cloud service type and explore various use cases for each cloud service type.

Learning objectives

After completing this module, you'll be able to:

- Describe infrastructure as a service (IaaS).
- Describe platform as a service (PaaS).
- Describe software as a service (SaaS).
- Identify appropriate use cases for each cloud service (IaaS, PaaS, SaaS).

2. Describe Infrastructure as a Service

<https://learn.microsoft.com/en-us/training/modules/describe-cloud-service-types/2-describe-infrastructure-service>

Describe Infrastructure as a Service


Completed

Infrastructure as a service (IaaS) is the most flexible category of cloud services. It provides the maximum amount of control for your cloud resources. In an IaaS model, the cloud provider is responsible for maintaining the hardware, network connectivity to the internet, and physical security. You're responsible for everything else, including:

- Operating system installation, configuration, and maintenance
- Network configuration
- Database and storage configuration

With IaaS, you're essentially renting the hardware in a cloud datacenter, but what you do with that hardware is up to you.

Earlier, you learned how the shared responsibility model divides duties between you and the cloud provider. The following diagram shows the individual infrastructure layers, such as networking, storage, servers, and runtime. It highlights which layers you operate in each service model.




Responsibility focus in IaaS

In IaaS, the cloud provider is responsible for the physical infrastructure and internet connectivity. You manage most of the workload stack, including operating systems, patching, configuration, and many security controls. This model gives you maximum flexibility and the largest operational responsibility.

Infrastructure as a Service (IaaS)

Maximum control — you manage the workload stack



Common Scenarios

Lift-and-shift migration

Set up cloud resources similar to on-premises, then move workloads to IaaS infrastructure.

Testing and development

Rapidly replicate dev/test environments. Start up or shut down with complete control.

Key takeaway

You rent hardware — what you do with it is up to you.

Scenarios

Common scenarios where IaaS might make sense include:

- **Lift-and-shift migration:** You set up cloud resources similar to your on-premises datacenter, and then move your workloads to the IaaS infrastructure.
- **Testing and development:** You need to rapidly replicate established configurations for development and test environments. You can start up or shut down different environments rapidly with an IaaS structure while maintaining complete control.

3. Describe Platform as a Service

<https://learn.microsoft.com/en-us/training/modules/describe-cloud-service-types/3-describe-platform-service>

Describe Platform as a Service


Completed

Platform as a service (PaaS) is a middle ground between renting space in a datacenter (infrastructure as a service) and paying for a complete and deployed solution (software as a service). In a PaaS environment, the cloud provider maintains the physical infrastructure, physical security, and connection to the internet. They also maintain the operating systems, middleware, development tools, and analytics services that make up a cloud solution. In a PaaS scenario, you don't have to worry about the licensing or patching for operating systems and databases.

PaaS is well suited to provide a complete development environment without the headache of maintaining all the development infrastructure.

Responsibility focus in PaaS

In PaaS, the cloud provider manages the physical infrastructure and platform components such as operating systems, middleware, and managed runtimes. You focus on your application code, data, and access controls. Depending on service configuration, some networking and application security settings are shared.



Scenarios

Common scenarios where PaaS might make sense include:

- **Development framework:** PaaS provides a framework that developers can build upon to develop or customize cloud-based applications. Developers can create applications using built-in software components. Cloud features such as scalability, high availability, and multitenant capability are included, reducing the amount of coding that developers must do.
- **Analytics or business intelligence:** Tools provided as a service with PaaS allow teams to analyze and mine their data, find insights and patterns, and predict outcomes to improve planning and operational decisions.

4. Describe Software as a Service

Describe Software as a Service


Completed

Software as a service (SaaS) is the most complete cloud service model from a product perspective. With SaaS, you're essentially renting or using a fully developed application. Email, financial software, messaging applications, and connectivity software are all common examples of a SaaS implementation.

While the SaaS model may be the least flexible, it's also the easiest to get up and running. It requires the least amount of technical knowledge or expertise to fully employ.

Responsibility focus in SaaS

In SaaS, the cloud provider manages almost all of the application stack, including infrastructure, platform, and application maintenance. You primarily manage your data, identity and access settings, and device access posture. SaaS has the lowest operational overhead for customers.



For example, a team using a SaaS collaboration platform can focus on user onboarding, access controls, and data governance while the provider handles infrastructure patching, application updates, and availability.

Scenarios

Common scenarios for SaaS include:

- Email and messaging.
- Productivity applications.
- Finance and expense tracking.

5. Module assessment

<https://learn.microsoft.com/en-us/training/modules/describe-cloud-service-types/5-knowledge-check>

Module assessment

Completed

6. Summary

<https://learn.microsoft.com/en-us/training/modules/describe-cloud-service-types/6-summary>

Summary

Completed

In this module, you learned about the cloud service types and some common scenarios for each type. You also reinforced how the shared responsibility model determines your responsibilities with different cloud service types.

Learning objectives

You should now be able to:

- Describe infrastructure as a service (IaaS).
- Describe platform as a service (PaaS).
- Describe software as a service (SaaS).
- Identify appropriate use cases for each cloud service (IaaS, PaaS, SaaS).

Explore with Copilot

Tip

Try one of these prompts in Copilot Chat:

- "For one application portfolio, choose IaaS, PaaS, or SaaS for each workload and explain the responsibility tradeoffs."
- "Show how operational responsibilities shift when moving from IaaS to PaaS, including security, patching, and scalability implications."
- "Create a short scenario quiz on cloud service types with answer explanations focused on selecting the right model."