Unit -2

Infrastructure as a Service (IaaS)

Infrastructure as a service (IaaS) is a form of cloud computing that provides virtualized computing resources over the internet. IaaS is one of the three main categories of cloud computing services, alongside software as a service (SaaS) and platform as a service (PaaS).

IaaS architecture and how it works

In an IaaS model, a cloud provider hosts the infrastructure components present in data center, including servers, storage and networking hardware,virtualization.

The IaaS provider also supplies a range of services to accompany those infrastructure components. These can include detailed billing, monitoring, log access, security, load balancing and clustering, as well as storage , such as backup, replication and recovery. These services are increasingly policy-driven, enabling IaaS users to implement greater levels of automation for important infrastructure tasks. For example, a user can implement policies to drive load balancing to maintain application availability and performance.

IaaS customers access resources and services through a wide area network (WAN), such as the internet, and can use the cloud provider's services to install the remaining elements of an application stack. For example, the user can log in to the IaaS platform to create virtual machines (VMs); install operating systems in each VM; deploy middleware, such as databases; create storage buckets for workloads and backups; and install the enterprise workload into that VM. Customers can then use the provider's services to track costs, monitor performance, balance network traffic, troubleshoot application issues, manage disaster recovery and more.

Any cloud computing model requires the participation of a provider. The provider is often a third-party organization that specializes in selling IaaS. Amazon Web services (AWS) and Google Cloud Platform (GCP) are examples of independent IaaS providers.

IaaS Key Features

- Highly scalable resources
- Enterprise-grade infrastructure
- Cost depends on consumption
- Multitenant architecture, i.e. a single piece of hardware serves many users
- The client gets complete control over the infrastructure

Benefits Of IaaS

IaaS offers many impressive benefits to the customers for ensuring affordability and for easily scaling the IT infrastructure. Benefits of implementing the IaaS environment include:

Pay Per Use

The IaaS service can be used on demand and the users only have to pay for the resources that are actually used.

Scalability

The IaaS infrastructure makes sure that the resources are available to the users when they need them. Therefore, there are no delays caused in the expansion of capacity and there is no wastage of the unused capacity.

Save Time And Cost

As the cloud service provider is responsible for setting up and maintaining the underlying physical hardware required for supporting the IaaS environment it saves a lot of time and effort of the users and ensures affordability.

Location Independence

Users working on the IaaS environment can access it from anywhere in the world through the internet; however, they have to abide by the security protocol of the cloud network.

Unaffected Service

There is no single point of failure in IaaS. Even though any one aspect of the hardware resources fail, the service will remain constant and unaffected.

Flexibility

One of the greatest benefits of IaaS include the ability to scale the resources up and down quickly according to the needs of the customers.

Faster Time To Market

Competition is an important factor in every business sector and faster time to market is one of the best ways to stay ahead of the competition. As the IaaS environment ensures flexibility and scalability, the business organizations can gear up and get their work done faster.

Focus On Business Growth

Business owners usually have to spend a lot of time, money and energy on making technology related decisions and recruiting staff for managing and maintaining their IT infrastructure. By opting for a service based IaaS model, business organizations can concentrate their time and resources where they are required.

Drawaback: Despite its flexible, pay-as-you-go model, IaaS billing can be a problem for some businesses. Cloud billing is extremely granular, and it is broken out to reflect the precise usage of services. It is common for users to experience sticker shock -- or finding costs to be higher than expected -- when reviewing the bills for every resource and service involved in an application deployment. Users should monitor their IaaS environments and bills closely to understand how IaaS is being used, and to avoid being charged for unauthorized services.

Insight is another common problem for IaaS users. Because IaaS providers own the infrastructure, the details of their infrastructure configuration and performance are rarely transparent to IaaS users. This lack of transparency can make systems management and monitoring more difficult for users.

IaaS users are also concerned about service resilience. The workload's availability and performance is highly dependent on the provider. If an IaaS provider experiences network bottlenecks or any form of internal or external downtime, the users' workloads will be affected. In addition, because IaaS is a multi-tenant architecture, the noisy neighbor issue can negatively impact users' workloads.

Major IaaS vendors and products

There are many examples of IaaS vendors and products. AWS offers storage services such as Simple Storage Services (S3) and Glacier, as well as compute services, including its Elastic Compute Cloud (EC2). GCP offers storage and compute services through Google Compute Engine (GCE), as does Microsoft Azure.

What are the biggest benefits, and challenges, you saw with IaaS adoption?

These are just a tiny sample of the broad range of services offered by major IaaS providers. Services can include serverless functions, such as AWS Lambda, Azure Functions or Google Cloud Functions; database access; big data compute environments; monitoring; logging; and more.

There are also many other smaller, or more niche players in the IaaS marketplace, including Rackspace Managed Cloud, CenturyLink Cloud, DigitalOcean and more.

Users will need to carefully consider the services, reliability and costs before choosing a provider -- and be ready to select an alternate provider and to redeploy to the alternate infrastructure if necessary.

How It Works?

In the IaaS technology, the cloud service provider hosts the IaaS infrastructure components that are traditionally present in a data center including network hardware, servers, storage and the virtualization of the hypervisor layer. The IaaS provider also provides a wide range of services to accompany the infrastructure components.

The Cloud

Just like all the cloud computing services, IaaS provides the users the access to computing resources in a virtualized environment. This is done through a public connection usually through the internet. IaaS provides the users the access to the virtualized environment for establishing their own IT platforms.

Virtualized Hardware

IaaS provides resources that are especially belonging to virtualized hardware which is also known as the computing infrastructure. The offerings in an IaaS environment include network connections, virtual server space, load balancers and IP addresses.

Cloud Servers

In physical terms, the cloud service provider extracts the pool of hardware resource from a group of servers and networks that are usually spread across various data centers and the cloud service provider is responsible for managing all the resources.

IaaS provider also offers relative services to the users for supporting the infrastructure components that include monitoring, detailed billing, security, load balancing, clustering along with storage services like data backup, replication and recovery. IaaS technology is focused at enabling the users to implement higher levels of automation for the crucial infrastructure tasks.

Examples Of IaaS

Internal Business Networks

Utilization of a pooled server and networking resources is done through which a business can store data and run applications. Growing businesses get the ability to scale their infrastructure according to the business growth.

Cloud Hosting

Hosting websites on virtual severs that are created on the pooled resources on the basis of the underlying physical servers.

Virtual Data Centers

A virtualized network is established that consists of virtual servers that can be used for offering advanced cloud hosting capabilities, enabling enterprise IT infrastructure or for integrating the operations.

Software as a Service (SaaS)

Software as a service (SaaS) is a software distribution model in which a third-party provider hosts applications and makes them available to customers over the Internet. SaaS is one of three main categories of cloud computing, alongside infrastructure as a service (IaaS) and platform as a service (PaaS).

SaaS is closely related to the application service provider (ASP) and on demand computing software delivery models. The hosted application management model of SaaS is similar to ASP, where the provider hosts the customer's software and delivers it to approved end users over the internet. In the software on demand SaaS model, the provider gives customers network-based access to a single copy of an application that the provider created specifically for SaaS distribution. The application's source code is the same for all customers and when new features or functionalities are rolled out, they are rolled out to all customers. Depending upon the service level agreement (SLA), the customer's data for each model may be stored locally, in the cloud or both locally and in the cloud.Organizations can integrate SaaS applications with other software using application programming interfaces (APIs). For example, a business can write its own software tools and use the SaaS provider's APIs to integrate those tools with the SaaS offering.

There are SaaS applications for fundamental business technologies, such as email, sales management, customer relationship management (CRM), financial management, human resource management (HRM), billing and collaboration. Leading SaaS providers include Salesforce, Oracle, SAP, Intuit and Microsoft.

SaaS applications are used by a range of IT professionals and business users, as well as C-level executives.

Advantages

SaaS removes the need for organizations to install and run applications on their own computers or in their own data centers. This eliminates the expense of hardware acquisition, provisioning and maintenance, as well as software licensing, installation and support. Other benefits of the SaaS model include:

Flexible payments: Rather than purchasing software to install, or additional hardware to support it, customers subscribe to a SaaS offering. Generally, they pay for this service on a monthly basis using a pay-as-you-go model. Transitioning costs to a recurring operating expense allows many businesses to exercise better and more predictable budgeting. Users can also terminate SaaS offerings at any time to stop those recurring costs.

Scalable usage: Cloud services like SaaS offer high vertical scalability, which gives customers the option to access more, or fewer, services or features on-demand.

Automatic updates: Rather than purchasing new software, customers can rely on a SaaS provider to automatically perform updates and patch management. This further reduces the burden on in-house IT staff.

Accessibility and persistence: Since SaaS applications are delivered over the Internet, users can access them from any Internet-enabled device and location.

Disadvantages of SaaS

Though using software as a service looks to be a very viable option for most of the businesses, there are some downsides too which need to be considered. We have listed some of the cons of SaaS development here -

1. Insufficient Data Security

This is one of the top concerns for companies who are looking to opt for a SaaS-based application model. Issues such as identity and access management need to be addressed before trusting any third party service provide with your company's sensitive data. Particularly in the case of accessibility from a mobile device, strict measures need to be taken before any kind of sensitive data is divulged to the service provider.

2. Difficulty with Regulations Compliance

   When your business critical data is stored in the service provider's data center, it is difficult to comply with the government's data protection regulations. Your company will need to learn which rules apply to your business, ask the right questions from your service provider, and address any kind of inconsistencies in the process.

3. Cumbersome Data Mobility

   The software as service market is filled with startups, and many of them do not have enough experience to survive in a highly competitive atmosphere. In case of a failure or in an event where you want to change your service provider, it becomes a cumbersome task to transfer your company's critical data from one service provider to another. Therefore, you need to be prepared for such an event with a planned exit strategy.

4. Low Performance

   A browser-based application running on a remote data center may lack in performance when compared to a similar application running from your employee's desktop. Companies therefore need to invest in a fast and reliable internet connection to negate this factor and also use tools for application performance management to know how their SaaS apps are performing over time.

5. Troublesome Software Integration

   When working with an external SaaS service provider to host multiple apps, there might be an integration problem with the existing in-house software. The in-house APIs and data structures might not integrate properly with the external software. As a result, you should always perform compatibility checks with all SaaS applications for better results.

But SaaS also poses some potential disadvantages. Businesses must rely on outside vendors to provide the software, keep that software up and running, track and report accurate billing and facilitate a secure environment for the business' data. Providers that experience service disruptions, impose unwanted changes to service offerings, experience a security breach or any other issue can have a profound effect on the customers' ability to use those SaaS offerings. As a result, users should understand their SaaS provider's service-level agreement, and make sure it is enforced.

Platform as a Service (PaaS)

Platform as a service (PaaS) is a cloud computing model in which a third-party provider delivers hardware and software tools -- usually those needed for application development -- to users over the internet. A PaaS provider hosts the hardware and software on its own infrastructure. As a result, PaaS frees users from having to install in-house hardware and software to develop or run a new application.

PaaS architecture and how it works

PaaS does not typically replace a business's entire IT infrastructure. Instead, a business relies on PaaS providers for key services, such as application hosting or Java development.

A PaaS provider builds and supplies a resilient and optimized environment on which users can install applications and data sets. Users can focus on creating and running applications rather than constructing and maintaining the underlying infrastructure and services.

Many PaaS products are geared toward software development. These platforms offer compute and storage infrastructure, as well as text editing, version management, compiling and testing services that help developers create new software more quickly and efficiently. A PaaS product can also enable development teams to collaborate and work together, regardless of their physical location.

PaaS pros and cons

The principal benefit of PaaS is simplicity and convenience for users -- the PaaS provider supplies much of the infrastructure and other IT services, which users can access anywhere via a web browser. PaaS providers then charge for that access on a per-use basis -- a model that many enterprises prefer, as it eliminates the capital expenses they traditionally have for on-premises hardware and software. Some PaaS providers charge a flat monthly fee to access their service, as well as the apps hosted within it.

Service availability or resilience, however, can be a concern with PaaS. If a provider experiences a service outage or other infrastructure disruption, this can adversely affect customers and result in costly lapses of productivity. Provider lock-in is another common concern, since users cannot easily migrate many of the services and much of the data produced through one PaaS product to another competing product. Users must evaluate the business risks of service downtime and lock-in before they commit to a PaaS provider.

Internal changes to a PaaS product are also a potential issue. For example, if a PaaS provider stops supporting a certain programming language or opts to use a different set of development tools, the impact on users can be difficult and disruptive. Users must follow the PaaS provider's service roadmap to understand how the provider's plans will affect its environment and capabilities.

PaaS vs. SaaS vs. IaaS

PaaS is one of three main categories of cloud computing services. The other two are software as a service (SaaS) and infrastructure as a service (IaaS).

With IaaS, a provider supplies the basic compute, storage and networking infrastructure along with the hypervisor (the virtualization layer). Users must then create virtual machines, install operating systems, support applications and data, and handle all of the configuration and management associated with those tasks.

With PaaS, a provider offers more of the application stack than IaaS providers, adding operating systems, middleware (such as databases) and other runtimes into the cloud environment.

With SaaS, a provider offers an entire application stack. Users simply log in and use the application that runs completely on the provider's infrastructure.

Leading PaaS vendors

There are many examples of PaaS providers that supply the tools and services needed to build enterprise applications in the cloud.

Google App Engine supports distributed web applications using Java, Python, PHP and Go. Red Hat OpenShift is a PaaS offering for creating open source applications using a wide variety of languages, databases and components. The Heroku PaaS offers Unix-style container computing instances that run processes in isolated environments, while supporting languages like Ruby, Python, Java.

While many PaaS providers offer similar services, each provider can pose unique nuances and limitations. It is important for users to test prospective providers to ensure their services meet any business or technical requirements, such as languages supported and service availability.

How does PaaS compare to internally hosted development environments?

PaaS can be accessed over any internet connection, making it possible to build an entire application in a web browser. Because the development environment is not hosted locally, developers can work on the application from anywhere in the world. This enables teams that are spread out across geographic locations to collaborate. It also means developers have less control over the development environment, though this comes with far less overhead.

What is included in PaaS?

The main offerings included by PaaS vendors are:

- Development tools
- Middleware
- Operating systems
- Database management
- Infrastructure

Different vendors may include other services as well, but these are the core PaaS services.

Development tools

PaaS vendors offer a variety of tools that are necessary for software development, including a source code editor, a debugger, a compiler, and other essential tools. These tools may be offered together as a framework. The specific tools offered will depend on the vendor, but PaaS offerings should include everything a developer needs to build their application.

Middleware

Platforms offered as a service usually include middleware, so that developers don't have to build it themselves. Middleware is software that sits in between user-facing applications and the machine's operating system; for example, middleware is what allows software to access input from the keyboard and mouse. Middleware is necessary for running an application, but end users don't interact with it.

Operating systems

A PaaS vendor will provide and maintain the operating system that developers work on and the application runs on.

Databases

PaaS providers administer and maintain databases. They will usually provide developers with a database management system as well.

Infrastructure

PaaS is the next layer up from IaaS in the cloud computing service model, and everything included in IaaS is also included in PaaS. A PaaS provider either manages servers, storage, and physical data centers, or purchases them from an IaaS provider.

Why do developers use PaaS?

Faster time to market

PaaS is used to build applications more quickly than would be possible if developers had to worry about building, configuring, and provisioning their own platforms and backend infrastructure. With PaaS, all they need to do is write the code and test the application, and the vendor handles the rest.

One environment from start to finish

PaaS permits developers to build, test, debug, deploy, host, and update their applications all in the same environment. This enables developers to be sure a web application will function properly as hosted before they release, and it simplifies the application development lifecycle.

Price

PaaS is more cost-effective than leveraging IaaS in many cases. Overhead is reduced because PaaS customers don't need to manage and provision virtual machines. In addition, some providers have a pay-as-you-go pricing structure, in which the vendor only charges for the computing resources used by the application, usually saving customers money. However, each vendor has a slightly different pricing structure, and some platform providers charge a flat fee per month.

Ease of licensing

PaaS providers handle all licensing for operating systems, development tools, and everything else included in their platform.

What are the potential drawbacks of using PaaS?

Vendor lock-in

It may become hard to switch PaaS providers, since the application is built using the vendor's tools and specifically for their platform. Each vendor may have different architecture requirements. Different vendors may not support the same languages, libraries, APIs, architecture, or operating system used to build and run the application. To switch vendors, developers may need to either rebuild or heavily alter their application.

Vendor dependency

The effort and resources involved in changing PaaS vendors may make companies more dependent on their current vendor. A small change in the vendor's internal processes or infrastructure could have a huge impact on the performance of an application designed to run efficiently on the old configuration. Additionally, if the vendor changes their pricing model, an application may suddenly become more expensive to operate.

Security and compliance challenges

In a PaaS architecture, the external vendor will store most or all of an application's data, along with hosting its code. In some cases the vendor may actually store the databases via a further third party, an IaaS provider. Though most PaaS vendors are large companies with strong security in place, this makes it difficult to fully assess and test the security measures protecting the application and its data. In addition, for companies that have to comply with strict data security regulations, verifying the compliance of additional external vendors will add more hurdles to going to market.

How is Platform-as-a-Service different from serverless computing?

PaaS and serverless computing are similar in that for both, all a developer has to worry about is writing and uploading code, and the vendor handles all backend processes. However, scaling is vastly different when using the two models. Applications built using serverless computing, or FaaS, will scale automatically, while PaaS applications will not unless programmed to do so. Startup times also vary greatly; serverless applications can be up and running almost instantly, but PaaS applications are more like traditional applications and have to be running most of the time or all of the time in order to be immediately available for users.

Another difference is that serverless vendors do not provide development tools or frameworks, as PaaS vendors do. And finally, pricing separates the two models. PaaS billing is not nearly as precise as in serverless computing, in which charges are broken down to the number of seconds or fractions of a second each instance of a function runs.

Summary:::

IaaS: Infrastructure as a Service

This is a virtual equivalent of a traditional data center. Cloud infrastructure providers use virtualization technology to deliver scalable compute resources such as server s, network s and storage to their clients. This is beneficial for the clients, as they don't have to buy personal hardware and manage its component s. Instead, they can deploy their platforms and application s within the provider's virtual machines that offer the same technologies and capabilities as a physical data center.

An IaaS provider is responsible for the entire infrastructure, but users have total control over it. In turn, users are responsible for installing and maintaining apps and operating systems, as well as for security, runtime, middleware and data.

IaaS users can compare the cost and performance of different providers in order to choose the best option, as they can access them through a single API.

IaaS Advantages

- The most flexible and dynamic model
- Cost-effective due to pay-as-you-go pricing
- Easy to use due to the automate d deployment of hardware
- Management tasks are virtualized, so employees have more free time for other tasks

IaaS Disadvantages

- Data security issues due to multitenant architecture
- Vendor outages make customers unable to access their data for a while

- The need for team training to learn how to manage new infrastructure

When to Use IaaS

IaaS can be especially advantageous in some situations:

- If you are a small company or a startup that has no budget for creating your own infrastructure
- If you are a rapidly growing company and your demands are unstable and changeable
- If you are a large company that wants to have effective control over infrastructure but pay only for the resources you actually use

Examples of IaaS

The best-known IaaS solution s vendors are Microsoft Azure, Google Compute Engine (GCE), Amazon Web Services (AWS), Cisco Metapod, DigitalOcean, Linode and Rackspace.

PaaS: Platform as a Service

PaaS in cloud computing is a framework for software creation delivered over the internet. This is the offering of a platform with built-in software components and tools, using which developer s can create, customize, test and launch applications. PaaS vendors manage servers, operating system updates, security patches and backups. Clients focus on app development and data without worrying about infrastructure, middleware and OS maintenance.

The main difference between IaaS and PaaS lies in the degree of control given to users.

PaaS Key Features

- Allows for developing, testing and hosting apps in the same environment
- Resources can be scaled up and down depending on business needs
- Multiple users can access the same app in development
- The user doesn’t have complete control over the infrastructure
- Web services and databases are integrated
- Remote teams can collaborate easily

PaaS Advantages

- PaaS-built software is highly scalable, available and multi-tenant, as it is cloud-based
- The development process is quickened and simplified
- Reduced expenses for creating, testing and launching apps
- Automated company policy
- Reduced amount of coding required
- Allows for easy migrating to the hybrid cloud

PaaS Disadvantages

- Data security issues
- Compatibility of existing infrastructure (not every element can be cloud-enabled)
- Dependency on vendor’s speed, reliability and support

When to Use PaaS

Such solutions are especially profitable to developers who want to spend more time coding, testing and deploying their applications. Utilizing PaaS is beneficial when:

- Multiple developers work on one project
- Other vendors must be included
- You want to create your own customized apps

Examples of PaaS

The best-known PaaS solutions vendors are Google App Engine, Amazon AWS, Windows Azure Cloud Services, Heroku, AWS Elastic Beanstalk, Apache Stratos and OpenShift.

SaaS: Software as a Service

With this offering, users get access to the vendor’s cloud-based software. Users don’t have to download and install SaaS applications on local devices, but sometimes they may need plugins. SaaS software resides on a remote cloud network and can be accessed through the web or APIs. Using such apps, customers can collaborate on projects, as well as store and analyze data.

SaaS is the most common category of cloud computing. The SaaS provider manages everything from hardware stability to app functioning. Clients are not responsible for anything in this model; they only use programs to complete their tasks. In this case, the client software experience is fully dependent on the provider.

SaaS Key Features

- The subscription model of utilizing
- No need to download, install or upgrade software
- Resources can be scaled depending on requirements
- Apps are accessible from any connected device
- The provider is responsible for everything

SaaS Advantages

- No hardware costs
- No initial setup costs
- Automated upgrades
- Cross-device compatibility
- Accessible from any location
- Pay-as-you-go model
- Scalability
- Easy customization

SaaS Disadvantages

- Loss of control
- Limited range of solutions
- Connectivity is a must

When to Use SaaS

Utilizing SaaS is most beneficial in the following situations:

- If your company needs to launch a ready-made software quickly
- For short-term projects that require collaboration
- If you use applications on a temporary basis
- For applications that need both web and mobile access

Examples of SaaS

The best-known SaaS solutions vendors are Google Apps, Dropbox, Gmail, Salesforce, Cisco WebEx, Concur, GoToMeeting, Office365.

The Difference Between IaaS, PaaS and SaaS

The table below provides a clear comparison of IaaS vs. PaaS vs. SaaS. Platform as a Service vs. Infrastructure as a Service gives less control to the user, but Platform as a Service vs. Software as a Service gives more control to the user. If you were to compare IaaS vs. SaaS, IaaS is the place you can move to and work from using available resources, while SaaS is a ready-made product you can utilize immediately without additional efforts.

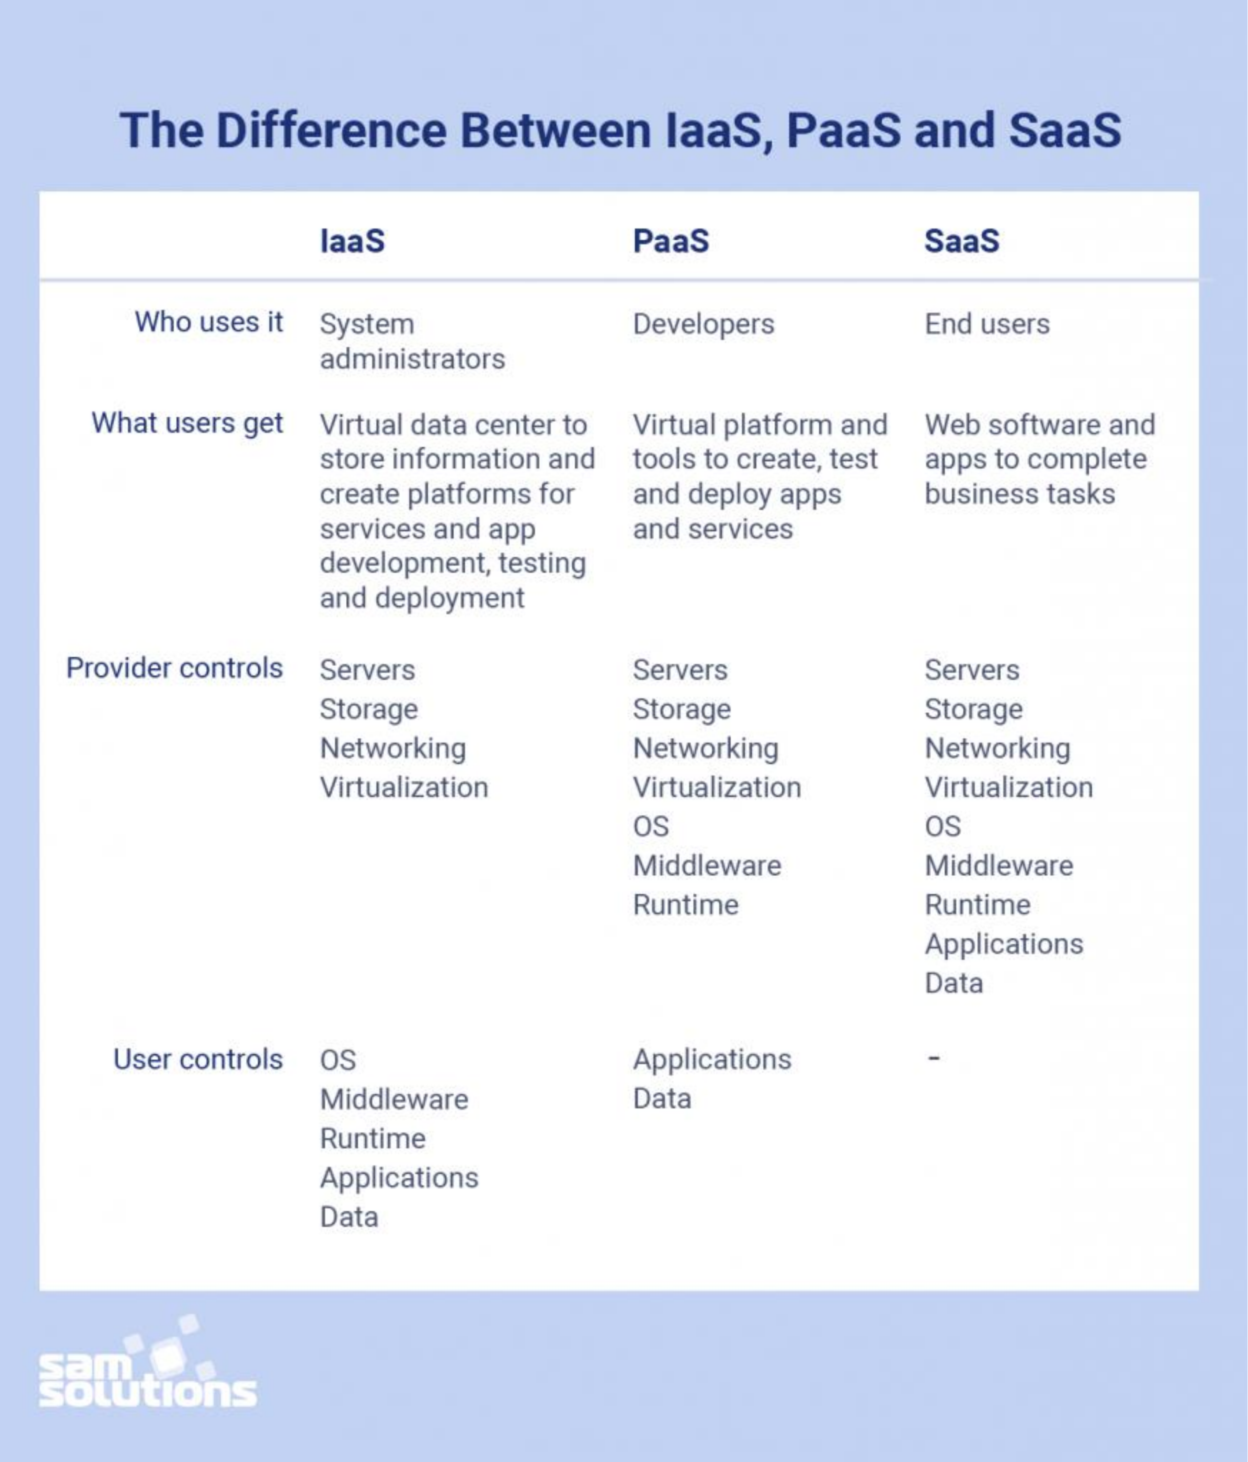

The Difference Between IaaS, PaaS and SaaS

	IaaS	PaaS	SaaS
Who uses it	System administrators	Developers	End users
What users get	Virtual data center to store information and create platforms for services and app development, testing and deployment	Virtual platform and tools to create, test and deploy apps and services	Web software and apps to complete business tasks
Provider controls	Servers Storage Networking Virtualization	Servers Storage Networking Virtualization OS Middleware Runtime	Servers Storage Networking Virtualization OS Middleware Runtime Applications Data
User controls	OS Middleware Runtime Applications Data	Applications Data	-

Unit -3

Service-level agreement (SLA)

A service-level agreement (SLA) is a contract between a service provider and its internal or external customers that documents what services the provider will furnish and defines the service standards the provider is obligated to meet.

Why are SLAs important?

Service providers need SLAs to help them manage customer expectations and define the circumstances under which they are not liable for outages or performance issues. Customers can also benefit from SLAs in that they describe the performance characteristics of the service, which can be compared with other vendors' SLAs, and also set forth the means for redressing service issues -- via service credits, for example.

For a service provider, the SLA is typically one of two foundational agreements it has with customers. Many service providers establish a master services agreement to establish the general terms and conditions in which they will work with customers. The SLA is often incorporated by reference into the service provider's master services agreement. Between the two service contracts, the SLA adds greater specificity regarding the services provided and the metrics that will be used to measure their performance.

What goes into an SLA?

In broad terms, an SLA will typically include a statement of objectives, a list of the services to be covered by the agreement and will also define the responsibilities of the service provider and customer under the SLA.

The customer, for example, will be responsible for making a representative available to resolve issues with the service provider in connection with the SLA. The service provider will be responsible for meeting the level of service as defined by the SLA. The service provider's performance is judged according to a set of metrics. Response time and resolution time are among the key metrics included in an SLA, since they relate to how the service provider deals with a service interruption.

A sampling of key features that may be included in an SLA

Performance metrics

SLAs establish customer expectations with regard to the service provider's performance and quality in a number of ways. Some metrics that SLAs may specify include the following:

- Availability and uptime percentage -- the amount of time services are running and accessible to the customer. Uptime is generally tracked and reported on per calendar month.
- Specific performance benchmarks to which actual performance will be periodically compared.
- Service provider response time -- the time it takes the service provider to respond to a customer's issue or request. A larger service provider may operate a service desk to respond to customer inquiries.
- Resolution time -- the time it takes for an issue to be resolved once logged by the service provider.
- The schedule for notification in advance of network changes that may affect users.
- Usage statistics that will be provided.

An SLA may specify availability, performance and other parameters for different types of customer infrastructure -- internal networks, servers and infrastructure components such as uninterruptable power supplies, for example.

Penalties: Repercussions for breaking terms

In addition to establishing performance metrics, an SLA may include a plan for addressing downtime and documentation for how the service provider will compensate customers in the event of a contract breach. Service credits are a typical remedy. Here, the service provider issues credits to the customer based on an SLA-specified calculation. Service providers, for example, might provide credits commensurate with the amount of time it exceeded the SLA's performance guarantee. A service provider may cap performance penalties at a maximum dollar amount to limit exposure.

The SLA will also include a section detailing exclusions, that is, situations in which an SLA's guarantees -- and penalties for failing to meet them -- don't apply. The list might include events such as natural disasters or terrorist acts. This section is sometimes referred to as a force majeure clause, which aims to excuse the service provider from events beyond its control.

Who needs a service-level agreement?

SLAs are thought to have originated with network service providers, but are now widely used in a range of IT-related fields. Companies that establish SLAs include IT service providers, managed service providers and cloud computing service providers. Corporate IT organizations, particularly those that have embraced IT service management (ITSM), enter SLAs with their in-house customers -- users in other departments within the enterprise. An IT department creates an SLA so that its services can be measured, justified and perhaps compared with those of outsourcing vendors.

Types of SLAs: Evolution

Over the years, SLAs have expanded to govern a growing set of IT procurement models. When IT outsourcing emerged in the late 1980s, SLAs evolved as a mechanism to govern such relationships. SLAs set the expectations for a service provider performance and established penalties for missing the targets and, in some cases, bonuses for exceeding them. Since outsourcing projects were frequently customized for a particular customer, outsourcing SLAs were often drafted to govern a specific project.

As managed services and cloud computing services became more prevalent in recent years, SLAs evolved to address those approaches. Shared services, rather than customized resources, characterize the newer contracting methods, so SLAs tend to be broad agreements intended to cover all of a service provider's customers.

SLAs, regardless of type, are subject to modification over time. Service providers will periodically review and update SLAs to reflect the addition of new services, changes to existing services or changes in the overarching regulatory environment.

What is a Service Level Agreement?

A Service Level Agreement (SLA) is a contract that describes the level of service a customer expects from his or her provider.

SLAs are used to establish measurable indicators of the service we provide thus ensure compliance with the expectations of our customers.

Service Desks such as ServiceTonic allow automating Service Level Agreements, making it easier to assign each ticket a priority or resolution time determined by the type of SLA.

Types of SLAs

3 types of SLA:

Service-based SLA

Applies a standard SLA to all customers that contract the same service. It is useful when our company offers several services with different resolution and response times.

For example, Premium or Standard services, incidence-kind or request-kind services, or any other distinction between services.

Customer-based SLA

Applies to all contracted services by a customer, a group of customers or the same business area.

For example, you can set a limit resolution time for incidents regarding “Budget” while prioritizing those that come from the “Finance Department” or an external customer.

Multilevel SLA

Combines service and customer SLA, and applies at a corporate level for all users in an organization too. Multilevel SLAs avoid duplication and incompetence between several agreements, making it possible to integrate several conditions into the same system.

For example, the Commercial Director may open requests by creating tickets that apply the standard SLA for the department, or a more restrictive SLA for “Business Management”, or an SLA for a specific service within the department such as “Suppliers”.

ServiceTonic helps you automate all SLA you want to define, even the most specific to your company.

Contact-based SLA

Applies to a specific user within a service that has a standard SLA. It is useful to offer a different treatment to a customer who we want to capture, retain or have a special attention to.

Service Level Management

Service Level Management is the process of managing service-level agreements. It is responsible for defining, documenting, agreeing, monitoring, measuring, reporting and reviewing the level of our services. It is what makes SLA a competitive advantage in our company.

Good SLA management lets us:

. Agree realistic conditions that our company can handle.
. Meet the expectations of our customers.
.Establish specific parameters for measuring the state of our services.
.Comply with the terms and conditions agreed with customers.
Avoid future conflicts. An agreement is a preventive communication to establish a transparent relationship. Therefore, meeting an SLA increases confidence.

How to design a good SLA

SLAs are a quality insurance that either contribute to customer loyalty or help us improve our services.

The most important condition when designing a good SLA is to ensure that our company is able to meet the agreement. To establish viable agreements, we will analyze both the service we provide and the internal structure we use to offer it.

Service Catalog

We will take a look at our Service Catalog for understanding the relationship between corporate areas involved and processes that are carried out when providing each service.

Operational Level Agreements

Consider automating internal SLAs to ensure compliance with customer expectations. These are called Operational Level Agreements, and are responsible for establishing an internal coordination to meet resolution and response times.

Surveys

SLA management does not end once we provide the service. It is important to analyze the level of satisfaction of our customers through, for example, periodic surveys. It is essential to ask for their opinion in order to know the effectiveness of our agreement management and adjust improvements if needed.

In short, SLA are essential for any service company. They strengthen the relationship with customers, who will clearly understand what we offer them while we know exactly what they expect from our company.

Use ServiceTonic to automate your SLAs and get real-time information in customizable dashboards, with notifications and alarms of your automated SLAs. Moreover, we will help you design your SLAs the most efficiently.

Unit-4

Virtualization is the "creation of a virtual (rather than actual) version of something, such as a server, a desktop, a storage device, an operating system or network resources".

The combination of hardware and software engineering that creates Virtual Machines (VMs) and enables multiple operating systems to run on the same platform. In the field of Information Technology, the fundamental change happening is Cloud Computing.

Virtualization is the backbone of Cloud Computing; Cloud Computing brings efficient benefits with the help of Virtualization.It also provides solutions for great challenges in the field of data security and privacy protection. Virtualization is the imitation of hardware within a software program. The role of multiple computers is allowed on a single computer. In a file or a web server, the usage of purchase, maintenance, depreciation, energy and floor space is double, but by creating virtual web or file server all of our objectives are fulfilled like the use of hardware resources to its maximum, flexibility, improvement in security, reduced cost. Efficient use of resources, increased security, portability, problem free testing, easier manageability, increased flexibility, fault isolation, rapid deployment are the benefits of virtualization.

Virtualization in Cloud Computing:

- For combining local and network resources data storage virtualization.
- For grouping physical storage devices into the single unit
- For reaching the high level of availability or improving availability using virtualization
- Improving performance using virtualization
- Using virtualization using stripping and caching
- Capacity improvement

A central computer hosting an application to multiple users, preventing the need for installing software repeatedly on every system **is** **virtualization in Cloud Computing**. The data from different hard drives, USB drives, and databases are merged into one location increasing its accessibility and security.

The creation of virtual hardware, software, or an operating system or a storage or network device is virtualization in cloud computing. In IT virtual changes occur more rapidly than in a physical environment. The changes occurring has to be managed, such changes are scalable and agile because of virtualization in Cloud Computing.

Importance of Virtualization:
For the maintenance of resources in cloud computing environment, virtualization is a necessity as it makes it easier. Virtualization in Cloud Computing increases security as it protects both the integrity of guest virtual machines and cloud components. Cloud Component virtualized machines can also be scaled up or down on demand or can provide reliability. Resource Sharing, high utilization of pooled resources, rapid provisioning are also some of the factors **Managed Service Provider VA** provides.

Reasons why you should use Managed Service Provider VA:

- Simplified Management
- Reduced system administrative work
- Resource Optimization
- It saves Money
- Easier software installation
- Data center consolidation and decreased power consumption
- Testing of CD's live without even burning them
- Better use from the hardware
- Increased CPU utilization
- Virtual machine can run on any x86 server

Virtualization is using computer resources to imitate other computer resources or whole computers. It separates resources and services from the underlying physical delivery environment.

Virtualization has three characteristics that make it ideal for cloud computing:

- **Partitioning:** In virtualization, many applications and operating systems (OSes) are supported in a single physical system by *partitioning* (separating) the available resources.
- **Isolation:** Each virtual machine is isolated from its host physical system and other virtualized machines. Because of this isolation, if one virtual-instance crashes, it doesn't affect the other virtual machines. In addition, data isn't shared between one virtual container and another.
- **Encapsulation:** A virtual machine can be represented (and even stored) as a single file, so you can identify it easily based on the service it provides. In essence, the encapsulated process could be a business service. This

encapsulated virtual machine can be presented to an application as a complete entity. Therefore, encapsulation can protect each application so that it doesn't interfere with another application.

Applications of virtualization

Virtualization can be applied broadly to just about everything that you could imagine:

- Memory
- Networks
- Storage
- Hardware
- Operating systems
- Applications

What makes virtualization so important for the cloud is that it decouples the software from the hardware. *Decoupling* means that software is put in a separate container so that it's isolated from operating systems.

Forms of virtualization

To understand how virtualization helps with cloud computing, you must understand its many forms. In essence, in all cases, a resource actually emulates or imitates another resource. Here are some examples:

- **Virtual memory:** Disks have a lot more space than computer memory. Therefore, with virtual memory, the computer frees valuable memory space by placing information it doesn't use often into disk space. PCs have *virtual memory,* which is a disk area that's used like memory. Although disks are very slow in comparison with memory, the user may never notice the difference, especially if the system does a good job of managing virtual memory. The substitution works surprisingly well.

- **Software:** Companies have built software that can emulate a whole computer. That way, one computer can perform as though it were actually 20 computers. The application consolidation results can be quite significant. For example, you might be able to move from a data center with thousands of servers to one that supports as few as a couple of hundred. This reduction results in less money spent not only on computers, but also on power, air conditioning, maintenance, and floor space.

- **Hypervisor**

A hypervisor or virtual machine monitor is computer software, firmware or hardware that creates and runs virtual machines. A computer on which a hypervisor runs one or more virtual machines is called a host machine, and each virtual machine is called a guest machine. Hypervisors provide several benefits to the enterprise data center. First, the ability of a physical host system to run multiple guest VMs can vastly improve the utilization of the underlying hardware. Where physical (nonvirtualized) servers might only host one operating system and application, a hypervisor virtualizes the server, allowing the system to host multiple VM instances -- each running an independent operating system and application -- on the same physical system using far more of the system's available compute resources.

Unit -5

Cloud Security

Cloud security, also known as cloud computing security, consists of a set of policies, controls, procedures and technologies that work together to protect cloud-based systems, data and infrastructure. These security measures are configured to protect data, support regulatory compliance and protect customers' privacy as well as setting authentication rules for individual users and devices. From authenticating access to filtering traffic, cloud security can be configured to the exact needs of the business. And because these rules can be configured and managed in one place, administration overheads are reduced and IT teams empowered to focus on other areas of the business.

The way cloud security is delivered will depend on the individual cloud provider or the cloud security solutions in place. However, implementation of cloud security processes should be a joint responsibility between the business owner and solution provider.

Why is Cloud Security Important?

For businesses making the transition to the cloud, robust cloud security is imperative. Security threats are constantly evolving and becoming more sophisticated, and cloud computing is no less at risk than an on-premise environment. For this reason, it is essential to work with a cloud provider that offers best-in-class security that has been customized for your infrastructure.

Cloud security offers many benefits, including:

Centralized security: Just as cloud computing centralizes applications and data, cloud security centralizes protection. Cloud-based business networks consist of numerous devices and endpoints. Managing these entities centrally enhances traffic analysis and filtering, streamlines the monitoring of network events and results in fewer software and policy updates. Disaster recovery plans can also be implemented and actioned easily when they are managed in one place.

Reduced costs: One of the benefits of utilizing cloud storage and security is that it eliminates the need to invest in dedicated hardware. Not only does this reduce capital expenditure, but it also reduces administrative overheads. Where once IT teams were firefighting security issues reactively, cloud security delivers proactive security features that offer protection 24/7 with little or no human intervention.

Reduced Administration: When you choose a reputable cloud services provider or cloud security platform, you can kiss goodbye to manual security configurations and almost constant security updates. These tasks can have a massive drain on resources, but when you move them to the cloud, all security administration happens in one place and is fully managed on your behalf.

Reliability: Cloud computing services offer the ultimate in dependability. With the right cloud security measures in place, users can safely access data and applications within the cloud no matter where they are or what device they are using.

More and more organizations are realizing the many business benefits of moving their systems to the cloud. Cloud computing allows organizations to operate at scale, reduce technology costs and use agile systems that give them the competitive edge. However, it is essential that organizations have complete confidence in their cloud computing security and that all data, systems and applications are protected from data theft, leakage, corruption and deletion.

All cloud models are susceptible to threats. IT departments are naturally cautious about moving mission-critical systems to the cloud and it is essential the right security provisions are in place, whether you are running a native cloud, hybrid or on-premise environment. Cloud security offers all the functionality of traditional IT security, and allows businesses to harness the many advantages of cloud computing while remaining secure and also ensure that data privacy and compliance requirements are met.

Choosing the Right Cloud Security Solution

Selecting the right cloud security solution for your business is imperative if you want to get the best from the cloud and ensure your organization is protected from unauthorized access, data breaches and other threats. **Forcepoint Cloud Access Security Broker (CASB)** is a complete cloud security solution that protects cloud apps and data, prevents compromised accounts and allows you to set security policies on a per-device basis. The result is a cloud infrastructure that is fully protected from known and emerging threats and which allows your organization to leverage the best that cloud computing has to offer.