

CLOUD COMPUTING

ELECTIVE – II

310254(C)

DR. (MISS) REVATI M WAHUL

***Course Objective 4:
To learn the application and
security on cloud computing***



Unit IV

Cloud Platforms and Cloud Applications

Points to cover

Part I : Amazon Web Services (AWS)

Amazon Web Services and Components

Amazon Simple DB

Elastic Cloud Computing (EC2)

Amazon Storage System

Amazon Database services (Dynamo DB)

Reading Material :

***Srinivasan, J. Suresh, "Cloud Computing: A Practical Approach for Learning and Implementation",
Chap 31, pg 336 – 343***

Points to cover

Part II : Microsoft Cloud Services

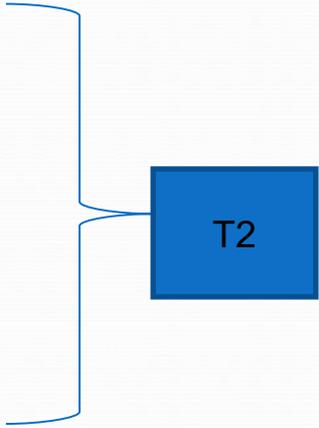
- Azure core concepts
- SQL Azure
- Windows Azure Platform Appliance.

Reading Material :

Rajkumar Buyya, Christian Vecchiola, S. Thamarai Selvi, "Mastering Cloud Computing Foundations and Applications Programming", Chap 9, pg 341– 350

Points to cover

Part III : Cloud Computing Applications

- Healthcare: ECG Analysis in the Cloud
 - Biology: Protein Structure Prediction
 - Geosciences: Satellite Image Processing
 - Business and Consumer Applications:
 - CRM and ERP
 - Social Networking
 - Google Cloud Application: Google App Engine. [T1]
 - Overview of OpenStack architecture. [Additional Material]
- 

Reading Material :

T2. Rajkumar Buyya, Christian Vecchiola, S. Thamarai Selvi, "Mastering Cloud Computing Foundations and Applications Programming", Chap 9, pg 353– 370

T1. Srinivasan, J. Suresh, "Cloud Computing: A Practical Approach for Learning and Implementation", Chap 30, pg 329 - 334

Amazon Web Services (AWS)

- Commenced in the year 2006
- A general cloud initiative
- Examples:
 - Amazon Elastic Compute Cloud (EC2)
 - Amazon Simple Storage Service (S3)
 - Amazon Aurora
 - Amazon DynamoDB
 - Amazon RDS
 - Amazon Lambda
 - Amazon VPC
 - Amazon Lightsail
 - Amazon SageMaker

1. Understanding Amazon Web Components And Services

- CC adopts scalable computing assets supplied as a service on pay-per-usage basis
- Any time, from any where over Internet
- User is unaware of the details
- CC is also utility computing or grid computing
- CC : a paradigm move in how we architect and consign scalable applications.

Earlier

- Heavy investment in construction infrastructure
- Approach:
 - ✓ Left large tracts of unused computing capability that took up space in large-scale data hubs
 - ✓ Needed somebody to babysit the servers
 - ✓ Had affiliated power costs
- The unused computing power trashed away, with no way to impel it out to other businesses or users who might be eager to pay for added compute cycles.

With CC, surplus computing capability can be put to use and be profitably traded to consumers.

Amazon Web Services (AWS)

Give programmatic entry to Amazon's ready-to-use computing infrastructure.

The robust computing podium that was assembled and enhanced by Amazon is now obtainable by anyone who has entry to the internet.

Amazon gives numerous web services, building-blocks that fulfill some of the quintessence requirements of most systems: storage, computing, messaging and datasets.

Amazon Web Services can aid us to architect scalable procedures by providing:

Reliability: The services run in Amazon's battle-tested, highly obtainable data centres that run Amazon's own business.

Security: Basic security and authentication methods are obtainable out of the packing box and customers can enhance them as wanted by layering his/her application-specific security on apex of the services.

Cost benefits: No fastened charges or support costs.

Ease of development: Simple APIs allow us to harness the full power of this virtual infrastructure and libraries, obtainable in most extensively employed programming languages.



Elasticity: Scale the computing supplies based on demand.

Cohesiveness: The four quintessence building-blocks using which services (storage, computing, messaging and datasets) are created from scratch currently work well and give a whole result through a large type of request for paid job domains.

Community: Tap into the vibrant and dynamic customer community that is propelling the extensive adoption of these web services and is bringing ahead sole requests for paid jobs assembled on this infrastructure.

AWS

There are two grades of supports accessible for users of Amazon Web Services:

1. Free forum-based support
2. Paid support packages

Amazon publishes the well-being rank of all its web services in a dashboard (accessible publicly) that is revised with any issues about the services.

Amazon presents standards-based SOAP and REST interfaces for combining with each of the services.

Developer libraries either from Amazon or third parties are accessible in multiple languages, encompassing Ruby, Python, Java, Erlang and PHP, for broadcasting with these services.

Amazon S3 (Storage): Simple Storage Service

Presents a web service interface for the storage and retrieval of data.

Any data can be retained and accessed from any location over the internet.

Users can store an unlimited number of things in S3, and the dimensions of each retained object can vary from 1 byte to 5 GB.

The data is retained securely utilizing the identical data storage infrastructure Amazon uses to power its worldwide mesh of e-commerce websites.

Limited access limits for each S3 object via HTTP requests.

Amazon S3 (Storage) Simple Storage Service

Absolutely relieves from its concern about storage space, get access to data or protecting the data.

Amazon double-checks high accessibility of the documents, so they are accessible when we require them.

The service-level affirmation supplied by Amazon for S3 commits to a 99.9% uptime, assessed monthly.

<https://aws.amazon.com/s3/pricing/>

Amazon EC2 (Elastic Computing)

A web service that permits us to use virtual machines within minutes and effortlessly scale the capability up or down founded on demand.

These instances are based on Linux and can run any submission or software.

The EC2 natural environment itself is constructed on the peak of the open source Xen hypervisor.

Amazon permits us to conceive Amazon Machine Images (AMIs) that act as templates for the instances.

Amazon EC2 (Elastic Computing)

Access to these can be controlled by identifying the permissions.

Presents factual web-scale computing, which makes it so straightforward to scale computing assets up and down.

Amazon presents five kinds of servers.

These servers vary from product single-core x86 servers to eight-core x86_64 servers.

AmazonSQS (Simple Queue Service)

Presents get access to the dependable messaging infrastructure utilized by Amazon.

Users can send and receive messages from any location utilizing direct REST-based HTTP requests.

The message is retained by Amazon over multiple servers and data hubs to supply the redundancy and reliability required from a messaging system.

Each message can comprise up to 8 KB of text data.

Amazon Simple Queue Service (SQS)

Incorporates very well with the other Amazon Web Services.

Presents a large way to construct a decoupled framework where EC2 instances can communicate with each other by dispatching a message to SQS and coordinating the workflow.

One can also use the queue for constructing a self-healing, auto-scaling EC2-based infrastructure for their application.

Amazon SimpleDB (SDB)

A web service for saving, processing and querying, organized datasets.

It is not a relational database in the customary sense, but it is a highly accessible schema, with a less-structured data shop in the cloud.

Simple to use and presents most of the purposes of a relational database.

The maintenance is much easier than a usual database.

Administrative tasks are taken care of by Amazon.

The data is mechanically indexed by Amazon and is accessible to us anytime from anywhere.

A key benefit of not being guarded to schemas is the proficiency to inject data on the fly and add new columns or keys dynamically.

2. Elastic Compute Cloud (EC2)

A most important component of Amazon.com's cloud computing platform, Amazon Web Services.

Permits scalable deployment of applications by supplying a web service.

Clients can use an AMI to conceive a virtual machine, encompassing any software programs desired.

A client can conceive, launch and terminate server instances as required, giving time for active servers, therefore the period 'elastic'.

Presents users with command over the geographical position of instances that permits for latency optimization and high grades of redundancy.

Amazon's features

A service grade affirmation for EC2

Microsoft Windows in beta pattern on EC2

Microsoft SQL Server in beta pattern on EC2

Designs for an AWS administration console

Designs for load balancing, auto-scaling and cloud supervising services

Amazon Elastic Compute Cloud (Amazon EC2) is a world wide web service that presents resizable computing capability that is utilized to construct and host software systems.

Major Constituents of EC2

1. Amazon Machine Images and Instances
2. Storage
3. Amazon S3
4. Databases
5. Amazon CloudWatch

Amazon Machine Images and Instances

A template that comprises a software program configuration (e.g., functioning scheme, submission server and applications).

From an AMI, a client can launch instances, which are running exact replicates of the AMI.

Also he/she can launch multiple instances of an AMI

Diagram

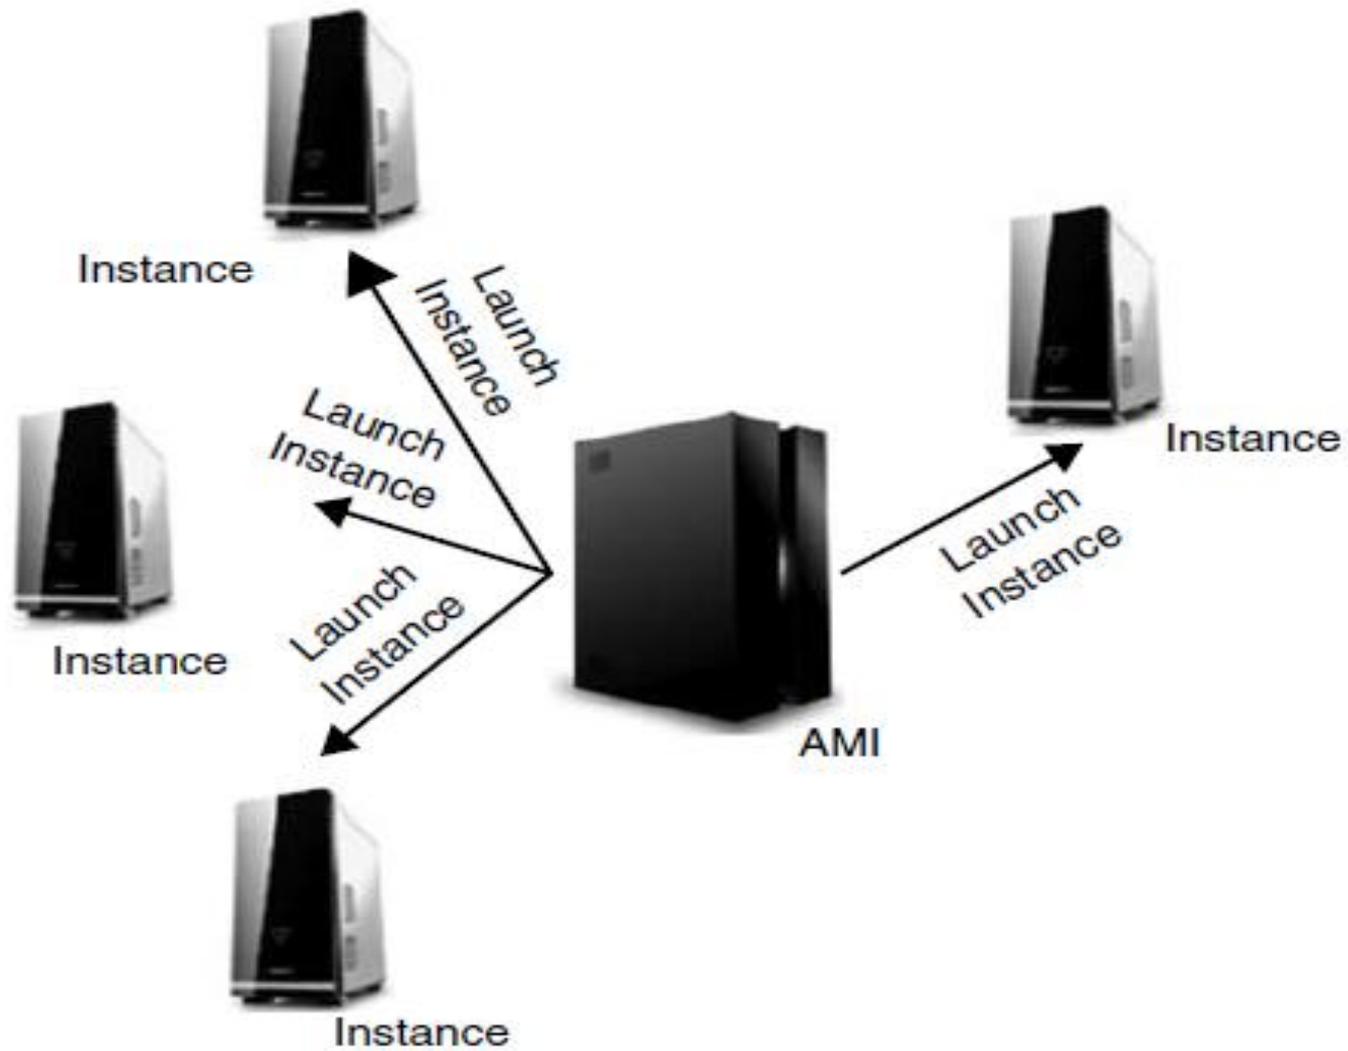


Figure 31.1 Amazon Machine Images and Instances

Amazon Machine Images and Instances

Instances hold running until the client halts or terminates them, or until they fail.

If an instance falls short, it can be commenced from the AMI.

One can use a single AMI or multiple AMIs counting on needs.

From a single AMI, distinct types of instances can be launched.

An instance type is vitally a hardware archetype.

Amazon publishes numerous AMIs that comprise widespread program configurations for public use.

In supplement, constituents of the AWS developer community have released their own made-to-order AMIs.

Storage

When utilizing EC2, data have to be stored.

The two most routinely utilized storage kinds are:

1. Amazon Simple Storage Service (Amazon S3)
2. Amazon Elastic Block Store (Amazon EBS) volumes

Amazon S3

storage for the internet

Presents a open World Wide Web service interface that provides us to shop and get any amount of data from any location on the web.

Amazon EBS presents the instances with continual, block-level storage.

Amazon EBS volumes are vitally hard computer disks that can adhere to a running instance.

Volumes are particularly matched for submissions that need a database, a document system or get access to raw block-level storage.

One can adhere multiple volumes to an instance.

To hold exact duplicate, a client can conceive a snapshot of the volume.

The user can also detach a volume from an instance and adhere it to a distinct one.

Databases

The submission running on EC2 might require a database.

Following are two well-known modes to apply a database for the application:

1. Use Amazon Relational Database Service (Amazon RDS), which provides us to effortlessly get an organized relational database in the cloud.
2. Launch an instance of a database AMI and use that EC2 instance as the database.

Amazon RDS boasts the benefit of managing the database administration jobs, for example, patching the programs, endorsing, saving backups, etc.

Amazon CloudWatch

A web service that presents real-time supervising to Amazon's EC2 clients on their asset utilization, such as CPU, computer disk and network.

Does not supply any recollection, computer disk space or load average metrics.

The data is aggregated and supplied through the AWS administration console.

It can additionally be accessed through online tools and web APIs.

The metrics assembled by Amazon CloudWatch provide auto scaling characteristics to dynamically add or eliminate EC2 instances.

The clients are recognized by the number of supervising instances.

3. Amazon Storage System S3

A service from Amazon permits the users to store documents into dependable isolated storage at a very comparable price.

Utilized by businesses to store photographs and videos of their clients, back up their own data and more.

Presents API for both SOAP and REST.

Manages objects and buckets.

An object agrees to a retained file.

Each object has an identifier, an owner and permissions.

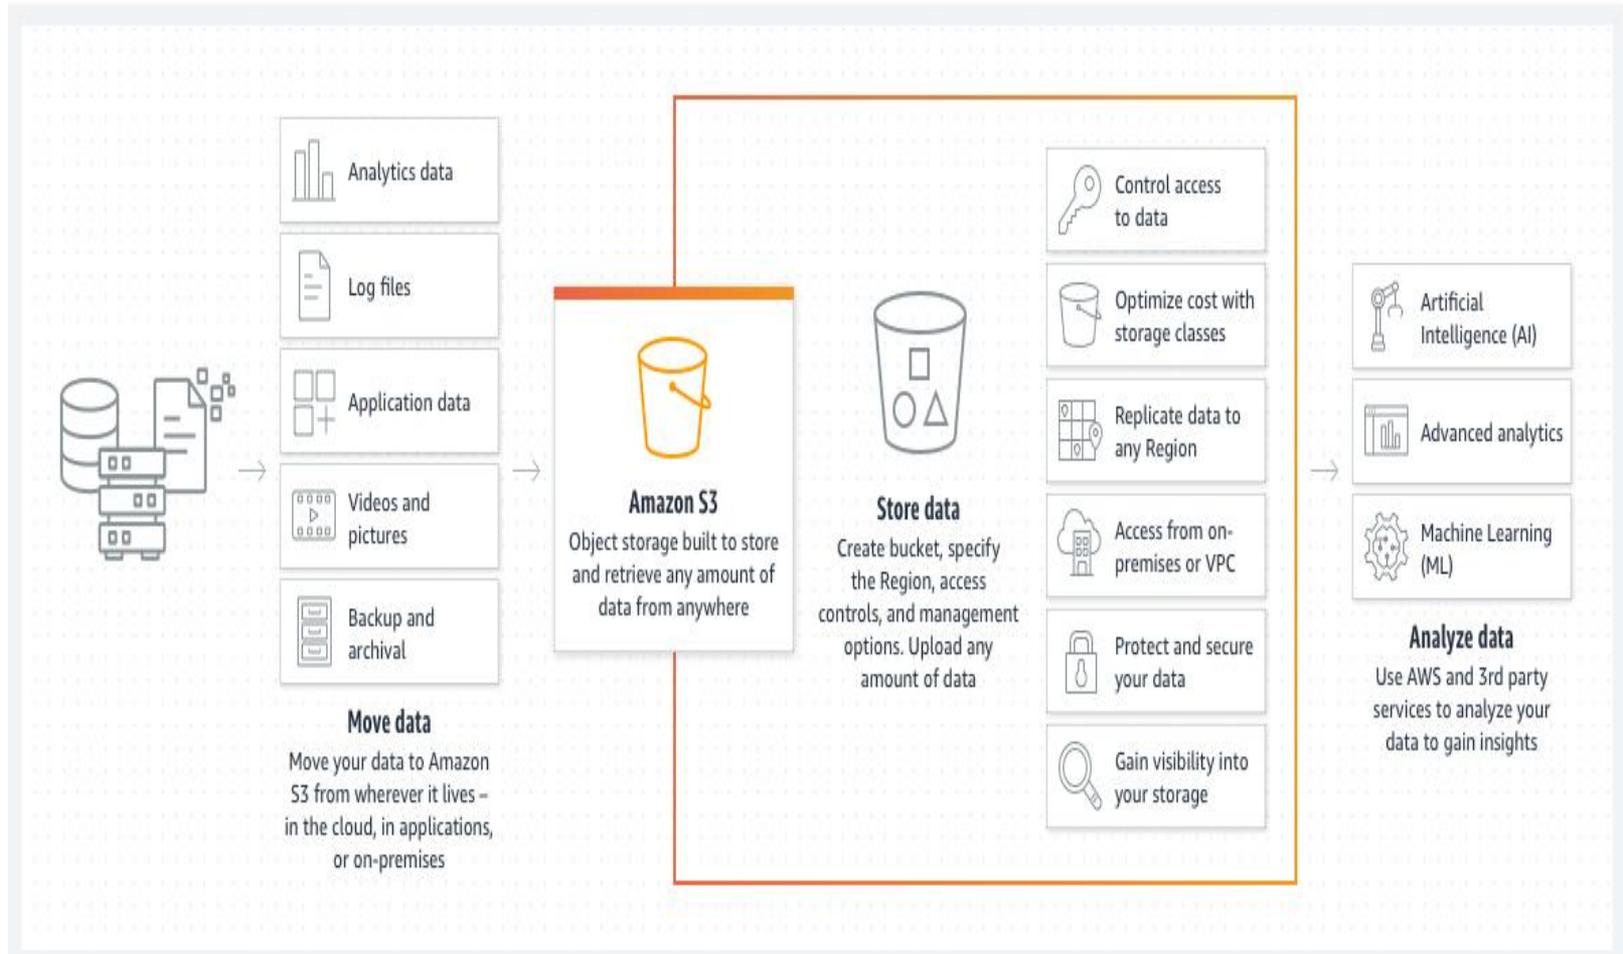
Objects are retained in a bucket.

A bucket has an exclusive name that should be well-trained with internet domain calling policy.

The object identifier consists of a filename with its relative path.

With this, S3 storage can emerge as a normal file system with folders and subfolders.

meet specific business, organizational, and compliance requirements.



S3

S3 stores random things (computer files) up to 5 terabytes in dimensions, each escorted by up to 2 kilobytes of metadata.

Objects are coordinated into buckets and recognized inside each bucket by an exclusive, user-assigned key.

Amazon Machine Images (AMIs) are changed in the Elastic Compute Cloud (EC2) that can be bundled to export to S3.

Amazon S3 is required for those who have the following issues:

Running out of bandwidths: Amazon S3 presents unlimited bandwidth and clients can be assisted with any amount of bandwidth the location needs.

Better scalability: Amazon S3 utilizes cloud hosting and assisting is somewhat fast. Separating them away from the usual HTTP demand will decisively lessen the server burden and therefore assures better stability.

Store documents online: Instead of endorsing up files in CDs/DVDs to save more computer hard disk space, another choice is to store them online. Users have the choice to hold documents with them in a personal account or make them publicly accessible.

Easier documents retrieval and sharing: If a document is online, the client can get access to it from any location as long as there is an internet connection. Amazon S3 furthermore permits one to communicate documents with other associates, purchasers and blog readers in a better way.

4. Amazon Database Services

a worldwide well-known web service

It makes relational database to set up, function and scale in the cloud much simpler.

It presents cost-efficient, resizable capability for multiple industry-standard relational databases and organizes widespread database management tasks.

Amazon RDS guides the users to get access to the capabilities of a well-renowned MySQL or Oracle database server.

Amazon RDS mechanically backs up the database and sustains the database programs that force the DB instance.

Amazon RDS is flexible.

Advantages of Amazon RDS

Accelerated deployment: decreases friction, when going from task development to deployment.

Can use direct **API calls** to get access to the capabilities of a production-ready relational database without being concerned about infrastructure provisioning or establishing and sustaining database software.

Managed: manages generic database administration tasks.

Compatible: One can easily get native access to a MySQL or Oracle database server with Amazon RDS.

Scalable: With a direct API call, clients can scale the compute and storage assets accessible to the database to meet enterprise desires and submission load.

Reliable: runs on the identical highly dependable infrastructure utilized by other AWS. Its automated backup service mechanically organizes the backup the database and lets us restore at any point.

Secure: presents World Wide Web service interfaces to configure firewall backgrounds that command network get access to the database instances and permits SSL (Secured Socket Layer) attachments to the DB Instances.

Inexpensive: In terms of deployment.

Amazon DynamoDB

Established on the values of Dynamo, a originator of NoSQL.

Adds the influence of the cloud to the NoSQL database world.

It claims the client's high availability, reliability and incremental scalability, with no restricts on dataset dimensions or demand throughput for a granted table.

Very fast

Runs on the newest in solid-state drive (SSD) expertise and integrates many other optimizations to commit reduced latency at any scale.



Amazon DynamoDB is a NoSQL database service that claims the benefits like:

- ✓ Managed
- ✓ Scalable
- ✓ Fast
- ✓ Durable
- ✓ Highly available
- ✓ Flexible

Points covered

Part I : Amazon Web Services (AWS)

- ✓ Amazon Web Services and Components
- ✓ Amazon Simple DB
- ✓ Elastic Cloud Computing (EC2)
- ✓ Amazon Storage System
- ✓ Amazon Database services (Dynamo DB)



Points to cover

Part II : Microsoft Cloud Services

Azure core concepts

SQL Azure

Windows Azure Platform Appliance

Microsoft Azure

A cloud operating system built on top of Microsoft datacenters' infrastructure

Provides developers with a collection of services for building applications with cloud technology

Services range from compute, storage, and networking to application connectivity, access control, and business intelligence.

Any application that is built on the Microsoft technology can be scaled using the Azure platform, which integrates the scalability features into the common Microsoft technologies such as Microsoft Windows Server 2008, SQL Server, and ASP.NET.

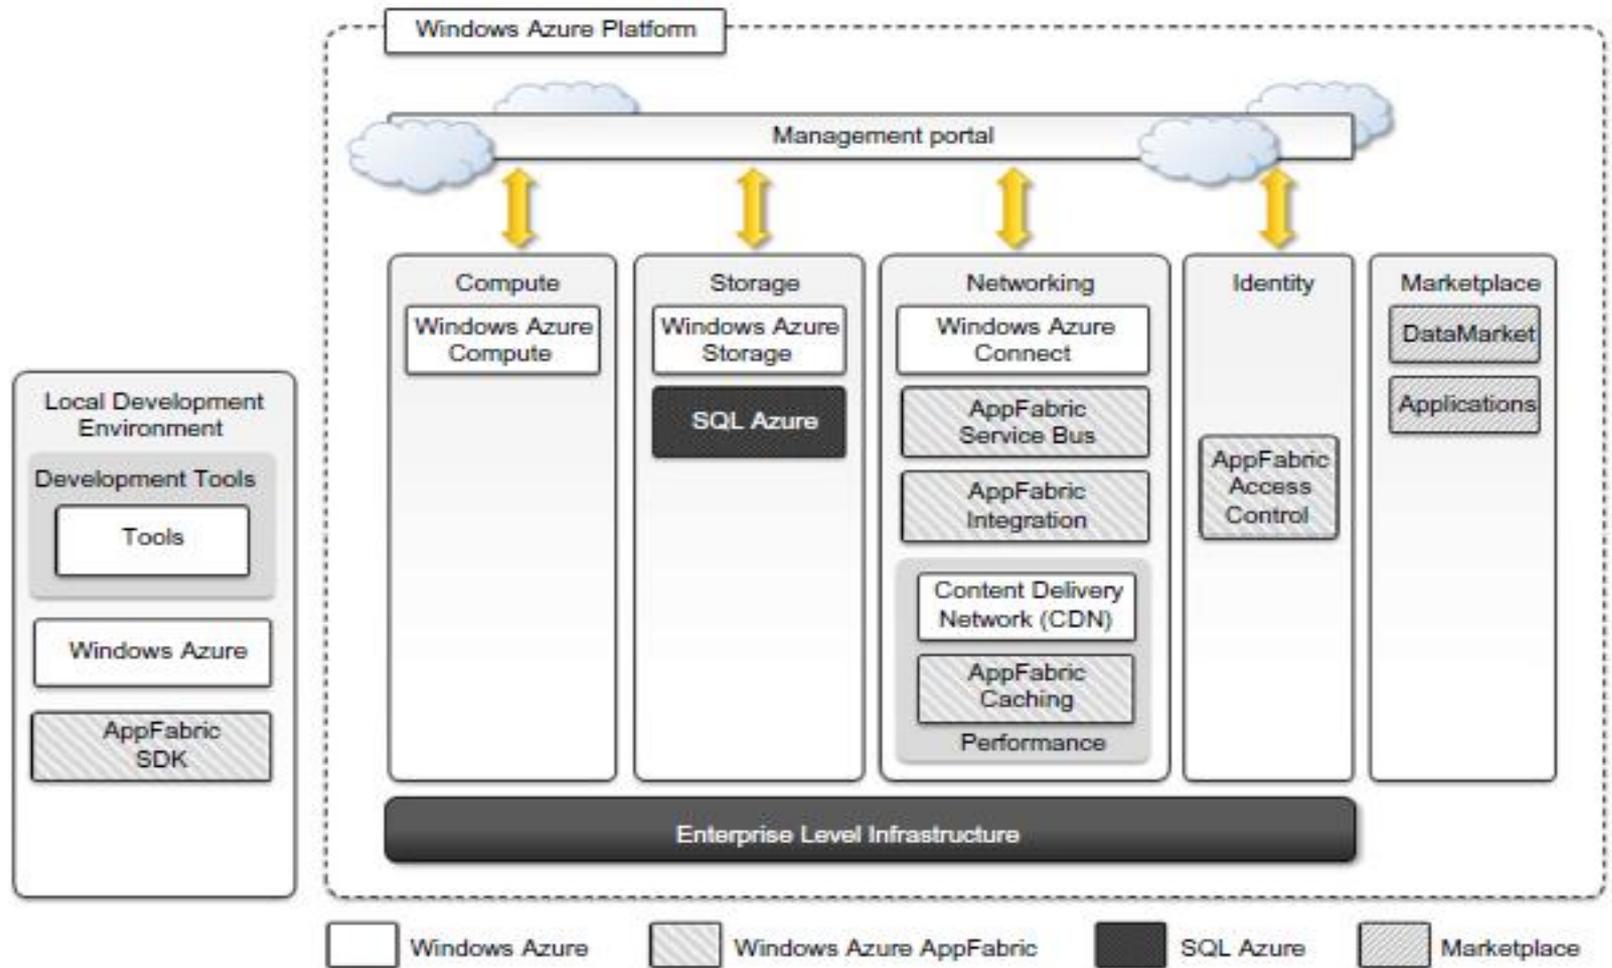


FIGURE 9.3

Microsoft Windows Azure Platform Architecture.

Microsoft Windows Azure Platform Architecture

Diagram: an overview of services provided by Azure
Services managed and controlled through the *Windows
Azure Management Portal*

- acts as an administrative console for all the services offered by the Azure platform

Azure core concepts

The Windows Azure platform is made up of

- a foundation layer and
- a set of developer services to be used to build scalable applications.

AppFabric : middleware that tie together services (compute, storage, networking, and identity management)

This scalable computing environment is hosted within Microsoft datacenters and accessible through the Windows Azure Management Portal.

Alternatively, developers can recreate a Windows Azure environment (with limited capabilities) on their own machines for development and testing purposes.

Overview of the Azure Middleware And its Services

1. Compute services
2. Storage services
3. Core infrastructure: AppFabric
4. Other services



1. Compute Services

Web role

Worker role

Virtual Machine role

Compute Services

The core components of Microsoft Windows Azure

Delivered by means of the abstraction of roles

Role:

- A runtime environment customized for a specific compute task
- Managed by the Azure operating system
- Instantiated on demand in order to address surges in application demand.
- 3 different roles:
 1. Web role
 2. Worker role
 3. Virtual Machine (VM) role

IIS 7 Web Server

a modular, extensible web server that was a complete redesign and rewrite of IIS, shipped with Windows Vista and Windows Server 2008, offering improved performance and security

Key Benefits:

Improved performance and reduced memory footprint.

Customizable installation options.

Support for HTML websites and optional support for ASP.NET, Classic ASP, and Web server extensions.

Ability to add or remove modules, and use protocols other than HTTP and HTTPS.

Web role

Designed to implement scalable Web applications

Represent the units of deployment of Web applications within the Azure infrastructure

Hosted on the IIS 7 Web Server, which is a component of the infrastructure that supports Azure

When Azure detects peak loads in the request made to a given application, it instantiates multiple Web roles for that application and distributes the load among them by means of a load balancer



Since version 3.5, the .NET technology natively supports Web roles

Developers can directly develop their applications in Visual Studio, test them locally, and upload to Azure.

It is possible to develop ASP.NET (ASP.NET Web Role and ASP.NET MVC 2 Web Role) and WCF (WCF Service Web Role) applications.

Since IIS 7 also supports the PHP runtime environment by means of the FastCGI module, Web roles can be used to run and scale PHP Web applications on Azure (CGI Web Role).

Other Web technologies that are not integrated with IIS can still be hosted on Azure (i.e., Java Server Pages on Apache Tomcat), but there is no advantage to using a Web role over a Worker role.

Worker role

Designed to host general compute services on Azure

Use to quickly provide compute power or to host services that do not communicate with the external world through HTTP.

Common Practice: To provide background processing for Web applications developed with Web roles.

Developing a worker role is like a developing a service.

Web role computation is triggered by the interaction with an HTTP client (i.e., a browser)

A Worker role runs continuously from the creation of its instance until it is shut down.

The Azure SDK provides developers with convenient APIs and libraries that allow connecting the role with the service provided by the runtime and easily controlling its startup as well as being notified of changes in the hosting environment.

As with Web roles, the .NET technology provides complete support for Worker roles, but any technology that runs on a Windows Server stack can be used to implement its core logic.

For example, Worker roles can be used to host Tomcat and serve JSP-based applications.

Virtual Machine Role

Allows developers to fully control the computing stack of their compute service by defining a custom image of the Windows Server 2008 R2 operating system and all the service stack required by their applications.

Based on the Windows Hyper-V virtualization technology which is natively integrated in the Windows server technology at the base of Azure.

Developers can image a Windows server installation complete with all the required applications and components, save it into a Virtual Hard Disk (VHD) file, and upload it to Windows Azure to create compute instances on demand.

Provides finer control of the compute service and resource that are deployed on the Azure Cloud compared to the Worker and Web roles.

An additional administrative effort is required for configuration, installation, and management of services.

Table 9.7 Windows Azure Compute Instances Characteristics, 2011–2012

Compute Instance Type	CPU	Memory	Instance Storage	I/O Performance	Hourly Cost (USD)
Extra Small	1.0 GHz	768 MB	20 GB	Low	\$0.04
Small	1.6 GHz	1.75 GB	225 GB	Moderate	\$0.12
Medium	2 × 1.6 GHz	3.5 GB	490 GB	High	\$0.24
Large	4 × 1.6 GHz	7 GB	1,000 GB	High	\$0.48
Extra Large	8 × 1.6 GHz	14 GB	2,040 GB	High	\$0.96

2. Azure Storage Services

Compute resources are equipped with local storage in the form of a directory on the local file system that can be used to temporarily store information that is useful for the current execution cycle of a role.

If the role is restarted and activated on a different physical machine, this information is lost.



Windows Azure provides different types of storage solutions that complement compute services with a more durable and redundant option compared to local storage.

Compared to local storage, these services can be accessed by multiple clients at the same time and from everywhere, thus becoming a general solution for storage.



Storage services

1. Blobs
2. Azure Drive
3. Tables
4. Queues

1. Blobs

Azure allows storing large amount of data in the form of binary large objects (BLOBs) by means of the blobs service.

This service is optimal to store large text or binary files.

Two types of blobs are available:

1. **Block blobs**
2. **Page blobs**

Blob Types

1. **Block blobs:**

- Composed of blocks and are optimized for sequential access
- Appropriate for media streaming
- Currently, blocks are of 4 MB, and a single block blob can reach 200 GB in dimension.

2. **Page blobs:**

- Made of pages that are identified by an offset from the beginning of the blob.
- Can be split into multiple pages or constituted of a single page.
- Optimized for random access and can be used to host data different from streaming.
- Currently, the maximum dimension of a page blob can be 1 TB.



Blobs storage provides users with the ability to describe the data by adding metadata.

It is also possible to take snapshots of a blob for backup purposes.

To optimize its distribution, blobs storage can leverage the Windows Azure CDN so that blobs are kept close to users requesting them and can be served efficiently.

2. Azure drive

Page blobs can be used to store an entire file system in the form of a single Virtual Hard Drive (VHD) file.

This can then be mounted as a part of the NTFS file system by Azure compute resources

To provide persistent and durable storage

A page blob mounted as part of an NTFS tree is called an Azure Drive.

3.Tables

Constitute a semi-structured storage solution, allowing users to store information in the form of entities with a collection of properties.

Entities are stored as rows in the table and are identified by a key, which also constitutes the unique index built for the table.

Users can insert, update, delete, and select a subset of the rows stored in the table.

Unlike SQL tables, there are no schema enforcing constraints on the properties of entities and there is no facility for representing relationships among entities.

Tables are more similar to spreadsheets rather than SQL tables.

Tables (2)

The service is designed to handle large amounts of data and queries returning huge result sets.

This capability is supported by partial result sets and table partitions.

A partial result set is returned together with a continuation token, allowing the client to resume the query for large result sets.

Table partitions allow tables to be divided among several servers for load-balancing purposes.

A partition is identified by a key, which is represented by three of the columns of the table.

Currently, a table can contain up to 100 TB of data, and rows can have up to 255 properties, with a maximum of 1 MB for each row.

The maximum dimension of a row key and partition keys is 1 KB

4. Queues

Queue storage allows applications to communicate by exchanging messages through durable queues

To avoid lost or unprocessed messages

Applications enter messages into a queue, and other applications can read them in a first-in, first-out (FIFO) style.

Queues (2)

To ensure that messages get processed, when an application reads a message it is marked as invisible

Hence, it will not be available to other clients.

Once the application has completed processing the message, it needs to explicitly delete the message from the queue.

This two-phase process ensures that messages get processed before they are removed from the queue, and the client failures do not prevent messages from being processed.

Queues (3)

At the same time, the queue does not enforce a strict FIFO model:

Messages that are read by applications that crash during processing are made available again after a timeout, during which other messages can be read by other clients.

An alternative to reading a message is peeking, which allows retrieving the message but letting it stay visible in the queue.

Messages that are peeked are not considered processed.



Note:

All the services are geo-replicated three times to ensure their availability in case of major disasters.

Geo-replication involves the copying of data into a different datacenter that is hundreds or thousands of miles away from the original datacenter.

Overview of the Azure Middleware And its Services

Compute services

Storage services

1. Core infrastructure: AppFabric
2. Other services

Core infrastructure: AppFabric

A comprehensive middleware for developing, deploying, and managing applications on the cloud or for integrating existing applications with cloud services.

Implements an optimized infrastructure supporting scaling out and high availability, sandboxing and multitenancy; state management; and dynamic address resolution and routing.

On top of this infrastructure, the middleware offers a collection of services that simplify many of the common tasks in a distributed application, such as communication, authentication and authorization, and data access.

These services are available through language-agnostic interfaces, thus allowing developers to build heterogeneous applications.

Core infrastructure: AppFabric

Access control

Service bus

Azure cache

<https://download.microsoft.com/download/c/a/7/ca7a9801-736e-421e-a386-16b37c83ffe8/windows-azure-appfabric-pdc10-overview.docx>

Access Control

AppFabric provides the capability of encoding access control to resources in Web applications and services into a set of rules that are expressed outside the application code base.

These rules give a great degree of flexibility in terms of the ability to secure components of the application and define access control policies for users and groups.

Access control services also integrate several authentication providers into a single coherent identity management framework.

Applications can leverage Active Directory, Windows Live, Google, Facebook, and other services to authenticate users.

This feature also allows easy building of hybrid systems, with some parts existing in the private premises and others deployed in the public cloud.

Service bus

Establishes the messaging and connectivity infrastructure provided with AppFabric for building distributed and disconnected applications in the Azure Cloud and between the private premises and the Azure Cloud.

Service Bus allows applications to interact with different protocols and patterns over a reliable communication channel that guarantees delivery.

Service Bus (2)

The service is designed

to allow transparent network traversal and

to simplify the development of loosely coupled applications,

without renouncing security and reliability and

letting developers focus on the logic of the interaction rather than the details of its implementation.

Allows services to be available by simple URLs, which are untied /freed from their deployment location.

It is possible to support publish-subscribe models, full-duplex communications point to point as well as in a peer-to-peer environment, unicast and multicast message delivery in one-way communications, and asynchronous messaging to decouple application components.

Service Bus (3)

In order to leverage these features, applications need to be connected to the bus, which provides these services.

A connection is the Service Bus element that is priced by Azure on a pay-as-you-go basis.

Users are billed :

- on a connections-per-month basis,
- they can buy advance “connection packs,” which have a discounted price, if they can estimate their needs in advance.



Azure cache

Windows Azure provides a set of durable storage solutions that allow applications to persist their data.

These solutions are based on disk storage, which might constitute a bottleneck for the applications that need to gracefully scale along the clients' requests and dataset size dimensions.

Azure Cache (2)

A service that allows developers to quickly access data persisted on Windows Azure storage or in SQL Azure.

The service implements a distributed in-memory cache
the cache size: dynamically adjusted by applications according to their needs

It is possible to store any .NET managed object as well as many common data formats (table rows, XML, and binary data) and control its access by applications.

Azure Cache is delivered as a service

Easily integrated with applications

This is particularly true for ASP.NET applications, which already integrate providers for session state and page output caching based on Azure Cache.

The service is priced according to the size of cache allocated by applications per month, despite their effective use of the cache.

Other services

Compute, storage, and middleware services constitute the core components of the Windows Azure platform.

Other services and components simplify the development and integration of applications with the Azure Cloud.

An important area for these services is applications connectivity, including virtual networking and content delivery.

Other Services:

1. Windows Azure virtual network
2. Windows Azure content delivery network

Windows Azure Virtual Network

Networking services for applications are offered under the name Windows Azure Virtual Network

Includes

1. Windows Azure Connect
2. Windows Azure Traffic Manager

1. Windows Azure Connect:

Allows easy setup of IP-based network connectivity among machines hosted on the private premises and the roles deployed on the Azure Cloud.

This service is particularly useful in the case of VM roles, where machines hosted in the Azure Cloud become part of the private network of the enterprise and can be managed with the same tools used in the private premises.

2. Windows Azure Traffic Manager

Provides load-balancing features for services listening to the HTTP or HTTPS ports and hosted on multiple roles.

It allows developers to choose from three different load-balancing strategies:

1. Performance,
2. Round-Robin, and
3. Failover.

Windows Azure Content Delivery Network

CDN is the content delivery network solution that improves the content delivery capabilities of Windows Azure Storage and several other Microsoft services, such as Microsoft Windows Update and Bing maps.

The service allows serving of Web objects (images, static HTML, CSS, and scripts) as well as streaming content by using a network of 24 locations distributed across the world.

Overview of the Azure Middleware And its Services

Compute services

Storage services

Core infrastructure: AppFabric

Other services

Points to cover

Part II : Microsoft Cloud Services

- ✓ Azure core concepts

 - SQL Azure

 - Windows Azure Platform Appliance

<https://learn.microsoft.com/en-us/azure/azure-sql/azure-sql-iaas-vs-paas-what-is-overview?view=azuresql>

SQL Azure

A relational database service hosted on Windows Azure and built on the SQL Server technologies

The service extends the capabilities of SQL Server to the cloud and provides developers with a scalable, highly available, and fault-tolerant relational database.

Accessible from either the Windows Azure Cloud or any other location that has access to the Azure Cloud.

Fully compatible with the interface exposed by SQL Server

Thus, applications built for SQL Server can transparently migrate to SQL Azure

The service is fully manageable using REST APIs

Allowing developers to control

- databases deployed in the Azure Cloud as well as
- the firewall rules set up for their accessibility

Diagram: the architecture of SQL Azure

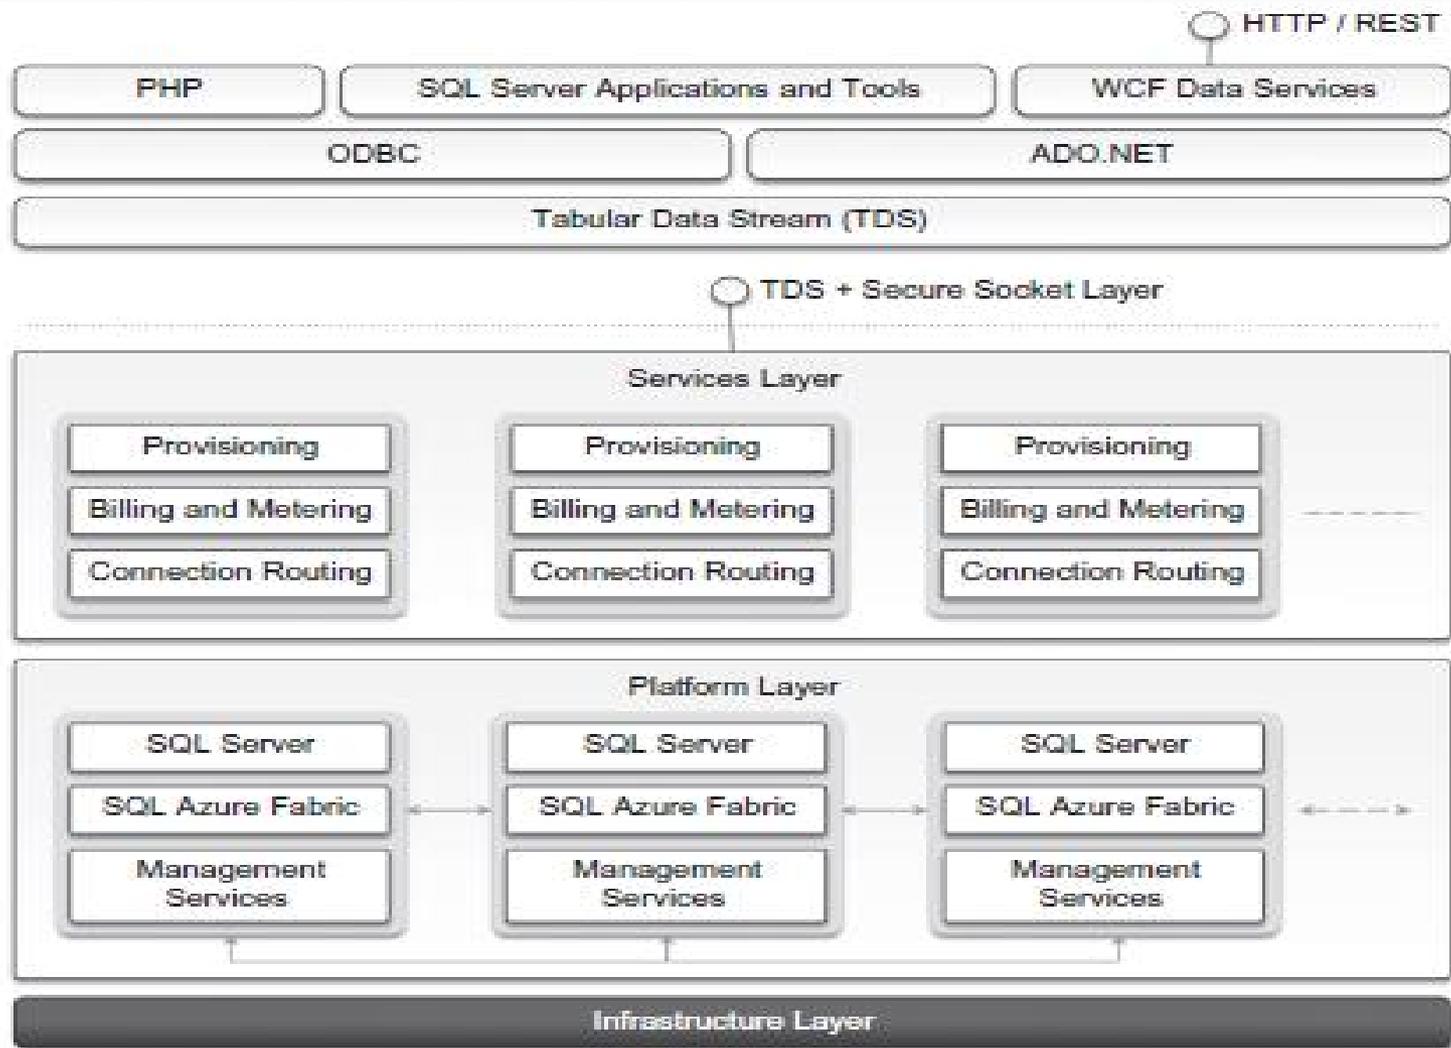


FIGURE 9.4

SQL Azure architecture.

SQL Azure

Access to SQL Azure is based on the TDS

Tabular Data Stream (TDS) protocol: the communication protocol underlying all the different interfaces used by applications to connect to a SQL Server-based installation such as ODBC and ADO.NET.

- *ADO.NET: ActiveX Data objects for the .NET Framework that provides communication between relational and non-relational systems through a common set of components.*

On the SQL Azure side, access to data is mediated by the service layer, which provides provisioning, billing, and connection-routing services.

These services are logically part of server instances, which are managed by SQL Azure Fabric.

This is the distributed database middleware that constitutes the infrastructure of SQL Azure and that is deployed on Microsoft datacenters.

Developers have to sign up for a Windows Azure account in order to use SQL Azure.

Once the account is activated, they can either use the Windows Azure Management Portal or the REST APIs to create servers and logins and to configure access to servers.

SQL Azure servers are abstractions that closely resemble physical SQL Servers:

They have a fully qualified domain name under the database.windows.net (i.e., server-name.database.windows.net) domain name.

This simplifies the management tasks and the interaction with SQL Azure from client applications.

SQL Azure ensures that multiple copies of each server are maintained within the Azure Cloud and that these copies are kept synchronized when client applications insert, update, and delete data on them.



The SQL Azure service is billed according to space usage and the type of edition.

Currently, two different editions are available: Web Edition and Business Edition.

1. The Web Edition is suited for small Web applications and supports databases with a maximum size of 1 GB or 5 GB.
2. The Business Edition is suited for independent software vendors, line-of-business applications, and enterprise applications and supports databases with a maximum size from 10 GB to 50 GB, in increments of 10 GB.

A bandwidth fee applies for any data transfer trespassing the Windows Azure Cloud or the region where the database is located.

A monthly fee per user/database is also charged and is based on the peak size the database reaches during the month.

Points to cover

Part II : Microsoft Cloud Services

- ✓ Azure core concepts

- ✓ SQL Azure

 - Windows Azure Platform Appliance

Windows Azure Platform Appliance

The Windows Azure platform can also be deployed as an appliance on third-party data centers and constitutes the cloud infrastructure governing the physical servers of the datacenter.

The Windows Azure Platform Appliance includes Windows Azure, SQL Azure, and Microsoft-specified configuration of network, storage, and server hardware.

The appliance is a solution that targets governments and service providers who want to have their own cloud computing infrastructure.

Windows Azure Platform Appliance(2)

Azure already provides a development environment that allows building applications for Azure in their own premises.

The local development environment is not intended to be production middleware, but it is designed for developing and testing the functionalities of applications that will eventually be deployed on Azure.

The Azure appliance is instead a full-featured implementation of Windows Azure.

Goal: to replicate Azure on a third-party infrastructure and make available its services beyond the boundaries of the Microsoft Cloud.

The appliance addresses two major scenarios:

1. institutions that have very large computing needs (such as government agencies) and
2. institutions that cannot afford to transfer their data outside their premises.

Points covered

Part I : Amazon Web Services (AWS)

- ✓ Amazon Web Services and Components
- ✓ Amazon Simple DB
- ✓ Elastic Cloud Computing (EC2)
- ✓ Amazon Storage System
- ✓ Amazon Database services (Dynamo DB)

Points covered

Part II : Microsoft Cloud Services

- ✓ Azure core concepts
- ✓ SQL Azure
- ✓ Windows Azure Platform Appliance

Points to cover

Part III : Cloud Computing Applications

Healthcare: ECG Analysis in the Cloud

Biology: Protein Structure Prediction

Geosciences: Satellite Image Processing

Business and Consumer Applications:

- CRM and ERP

- Social Networking

Google Cloud Application: Google App Engine.

Overview of OpenStack architecture.



Use Cases of Cloud Applications

Scientific Applications

Health care Analysis in the Cloud

Biology

Geo Science

Business and Consumer Applications

Productivity

Social Networking

Media Applications

Multiplayer online gaming

Cloud Applications

1. Scientific applications

Healthcare: ECG analysis in the cloud

Biology: protein structure prediction

Geoscience: satellite image processing

Healthcare : ECG Analysis in the Cloud

Healthcare is a domain where computer technology has found several and diverse applications: from supporting the business functions to assisting scientists in developing solutions to cure diseases.

An important application is the use of Cloud technologies for supporting doctors in providing more effective diagnostic processes.

electrocardiogram (ECG) data analysis on the Cloud

Healthcare : ECG Analysis in the Cloud

The development of Internet connectivity and its accessibility from any device at any time has made Cloud technologies an attractive option for developing health-monitoring systems.

Electrocardiogram (ECG) data analysis and monitoring constitutes a case study that naturally fits in this scenario.

ECG is the electrical manifestation of the contractile activity of the heart's myocardium.

This activity produces a specific waveform that is repeated overtime and that represents the heartbeat.

The analysis of the shape of the waveform is used to identify arrhythmias and it is the most common way for detecting heart diseases.

Healthcare : ECG Analysis in the Cloud (2)

Cloud computing technologies allow the remote monitoring of a patient's heartbeat data, its analysis in minimum time, and the notification of first-aid personnel and doctors should this data reveal potentially dangerous conditions.

This way a patient at risk can be constantly monitored without going to hospital for ECG analysis.

At the same time, doctors and first-aid personnel can instantly be notified with cases that require their attention.

Diagram : An online health monitoring system hosted in the cloud

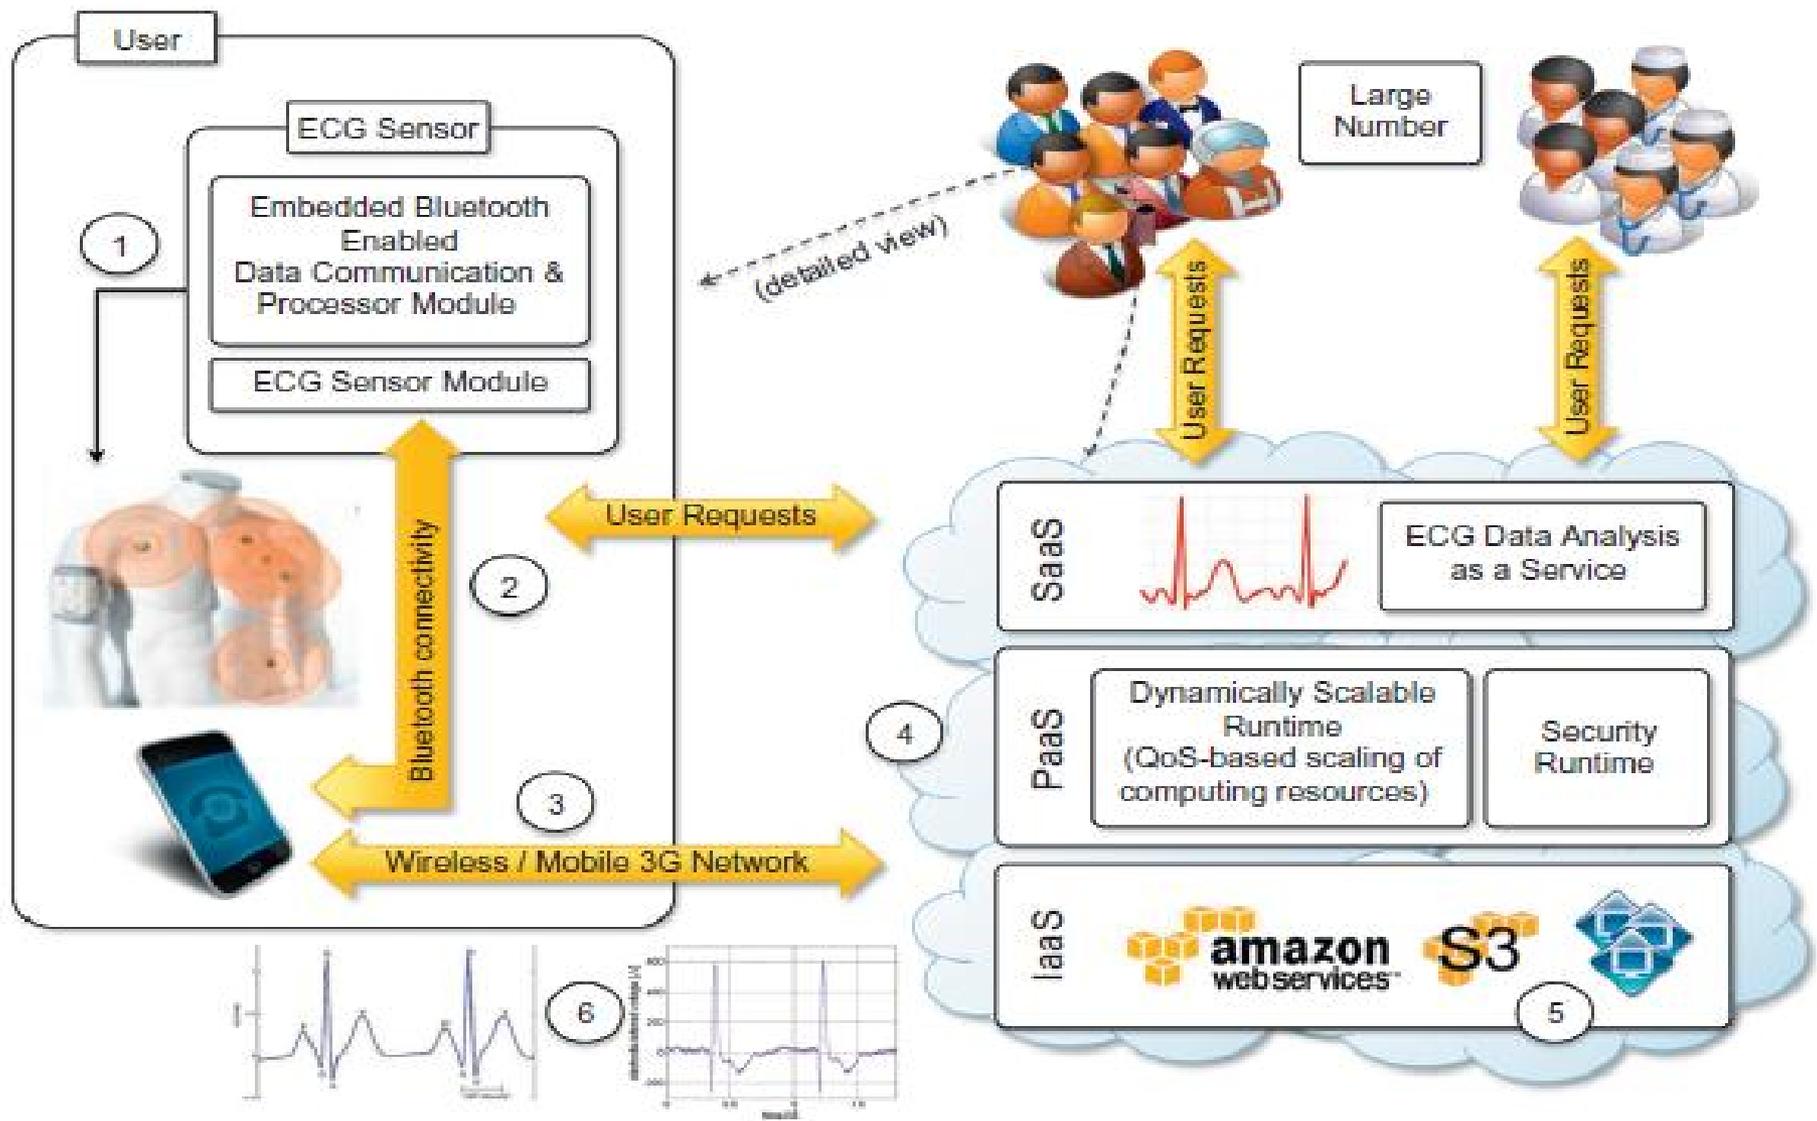


FIGURE 10.1

An online health monitoring system hosted in the cloud.



Wearable computing devices equipped with ECG sensors constantly monitor the patient's heartbeat.

Such information is transmitted to the patient's mobile device, which will eventually forward it to the cloud-hosted Web service for analysis.

The Web service forms the front-end of a platform that is entirely hosted in the cloud and that leverages the three layers of the cloud computing stack: SaaS, PaaS, and IaaS.

The Web service constitute the SaaS application that will store ECG data in the Amazon S3 service and issue a processing request to the scalable cloud platform.



The runtime platform is composed of a dynamically sizable number of instances running the workflow engine and Aneka [PaaS Platform].

The number of workflow engine instances is controlled according to the number of requests in the queue of each instance, while Aneka controls the number of EC2 instances used to execute the single tasks defined by the workflow engine for a single ECG processing job.

Each of these jobs consists of a set of operations involving the extraction of the waveform from the heartbeat data and the comparison of the waveform with a reference waveform to detect anomalies.

If anomalies are found, doctors and first-aid personnel can be notified to act on a specific patient.

Healthcare : ECG Analysis in the Cloud

Even though remote ECG monitoring does not necessarily require Cloud technologies, CC introduces opportunities that would be otherwise hardly achievable.

The first advantage is the **elasticity of the Cloud infrastructure** that can grow and shrink according to the requests served. As a result, doctors and hospitals do not have to invest in large computing infrastructures designed after capacity planning, thus making a more effective use of budgets. The second advantage is **ubiquity**.

Cloud computing technologies have now become easily accessible and promise to deliver systems with minimum or no downtime. Computing systems hosted in the Cloud are accessible from any Internet device through simple interfaces (such as SOAP and REST based web services).

This makes not only these systems ubiquitous but they can also be easily integrated with other systems maintained in the hospital's premises. Lastly, **cost savings** constitute another reason.

Cloud services are priced on a pay-per-use basis and with volume prices in case of large numbers of service requests. These two models provide a set of flexible options that can be used to price the service, thus actually charging costs based on effective use rather than capital costs.



Biology: protein structure prediction

Home assignment



Geoscience: satellite image processing

Home Assignment

Business and Consumer applications

Capex to Opex

Ubiquity

Elasticity

The combination of all these elements has made cloud computing the preferred technology for a wide range of applications, from CRM and ERP systems to productivity and social-networking applications.

CRM and ERP

Customer relationship management (CRM) and enterprise resource planning (ERP) applications are market segments that are flourishing in the cloud, with CRM applications the more mature of the two.

Cloud CRM applications create a great opportunity for small enterprises and start-ups to have fully functional CRM software without large up-front costs and by paying subscriptions.

Moreover, CRM is not an activity that requires specific needs, and it can be easily moved to the cloud.

Such a characteristic, together with the possibility of having access to your business and customer data from everywhere and from any device, has fostered the spread of cloud CRM applications.



CRM and ERP

ERP solutions on the cloud are less mature and have to compete with well-established in-house solutions.

ERP systems integrate several aspects of an enterprise: finance and accounting, human resources, manufacturing, supply chain management, project management, and CRM.

Goal : to provide a uniform view and access to all operations that need to be performed to sustain a complex organization.

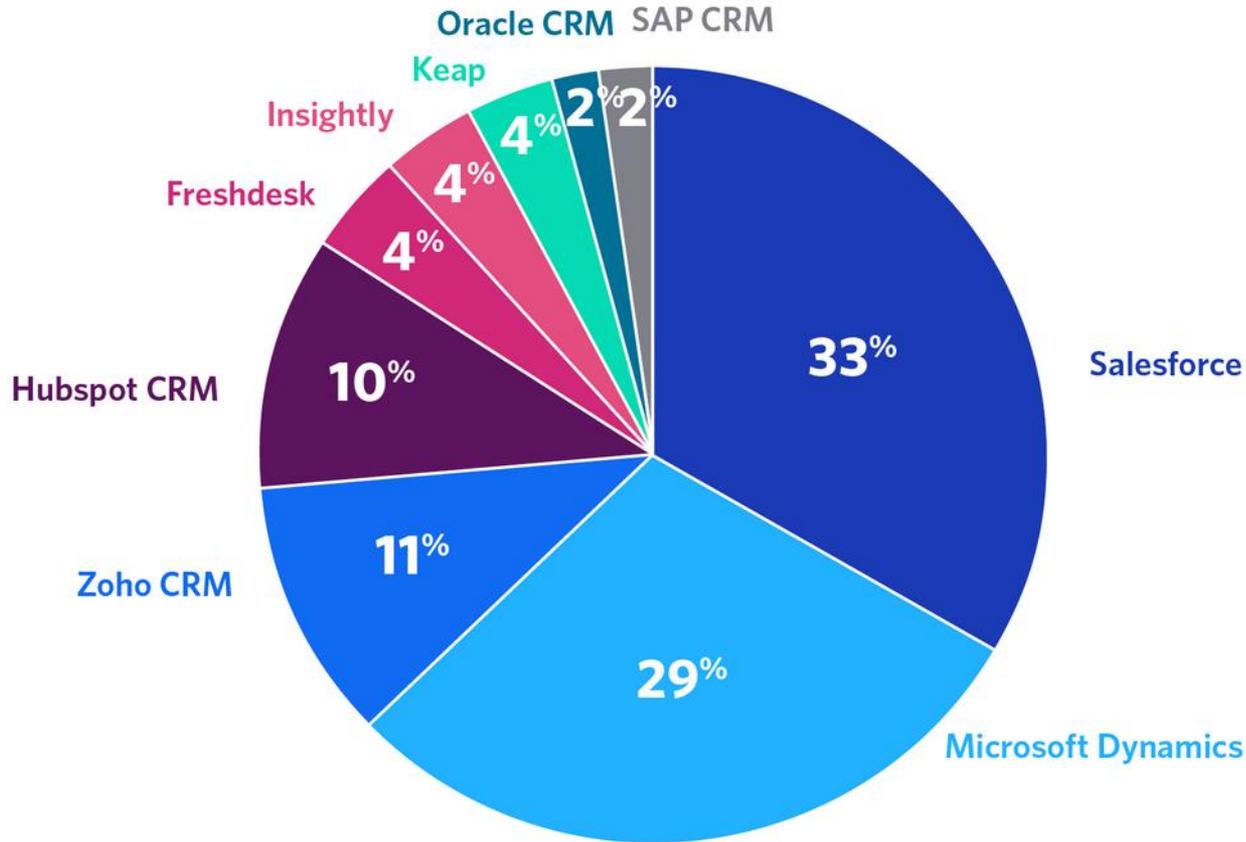
Because of the organizations that they target, the transition to cloud-based models is more difficult: the cost advantage over the long term might not be clear, and the switch to the cloud could be difficult if organizations already have large ERP installations.

For this reason cloud ERP solutions are less popular than CRM solutions at this time.

Top CRM Softwares



Market Share of Top CRM Software in 2021



Source: TrustRadius platform data collected in June 2021



CRM PLATFORMS

MARKET MAP

SOFTWARE LIST

BEST ERP SOFTWARE IN 2023

Download our 16 software picks to transform organizational efficiency

ORACLE NETSUITE

NetSuite ERP

☰ SUMMARY

A cloud-based ERP system that provides integrated solutions for finance, inventory management, order management, and customer relationship management.

▲ CUSTOMER REFERENCES

Williams-Sonoma, GoPro

🔍 EN SAVOIR PLUS

[Check website](#)

🕒 SOFTWARE TYPE

ERP



Infor M3

☰ SUMMARY

A cloud-based, manufacturing and distribution ERP system with new features to improve user productivity. Key modules include Supply Chain Management, ...

▲ CUSTOMER REFERENCES

Triumph, Burton, Flatiron, Belmond

🔍 EN SAVOIR PLUS

[Check website](#)

🕒 SOFTWARE TYPE

ERP



Sage X3

☰ SUMMARY

Sage is committed to small and medium enterprises worldwide. It is a market leader for integrated tools including accounting, payroll and payment system...

▲ CUSTOMER REFERENCES

RIVR Media, Tree Tops Caravan Park, BKP...

🔍 EN SAVOIR PLUS

[Check website](#)

🕒 SOFTWARE TYPE

ERP



Microsoft Dynamics 365

☰ SUMMARY

Manage the following enterprise functions with Dynamics 365, Sales and marketing, Service, Finance, Supply chain and more. Integrate the ERP with any Microsoft ...

▲ CUSTOMER REFERENCES

CocaCola, HP, Chevron, Sandvik, L'Oreal

🔍 EN SAVOIR PLUS

[Check website](#)

🕒 SOFTWARE TYPE

ERP

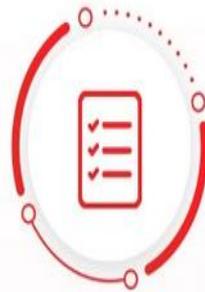
TOP 5 ERP MODULE



Financial
Management
Module



Human
Resources
Module



Order
Management
Module



Supply Chain
Management
Module



Warehouse
Management
Module

Top ERP Systems by Market Share



ORACLE

ENTERPRISE
RESOURCE PLANNING
CLOUD



SAP® Business
One



sage 100cloud

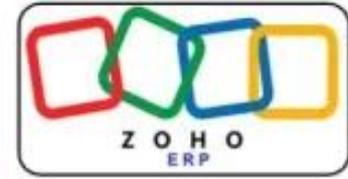


SAP S/4 HANA



infor

Popular ERP Software in India



Salesforce.com

Probably the most popular and developed CRM solution available today

More than 100,000 customers have chosen Salesforce.com to implement their CRM solutions.

The application provides customizable CRM solutions that can be integrated with additional features developed by third parties.

Salesforce.com is based on the Force.com cloud development platform.

This represents scalable and high-performance middleware executing all the operations of all Salesforce.com applications.

Diagram : The architecture of the Force.com platform

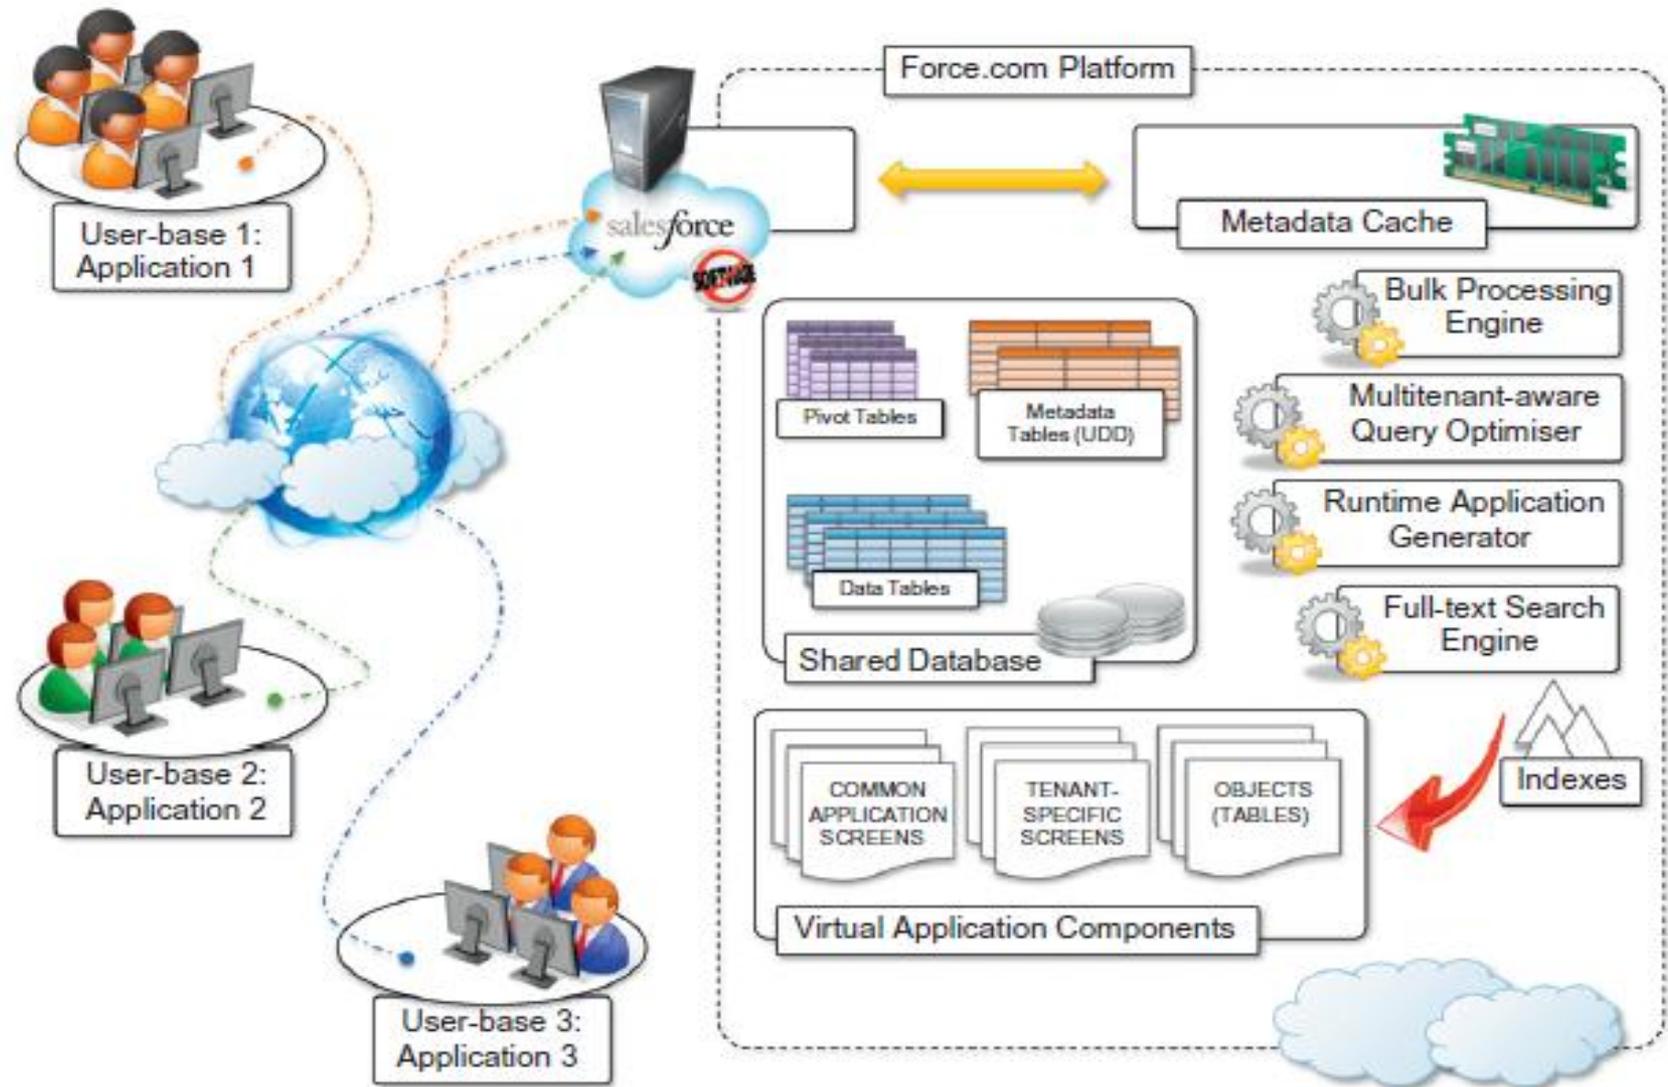


FIGURE 10.5

Salesforce.com and Force.com architecture.

Force.com

Initially designed to support scalable CRM applications, the platform has evolved to support the entire life cycle of a wider range of cloud applications by implementing a flexible and scalable infrastructure.

At the core of the platform resides its metadata architecture, which provides the system with flexibility and scalability.

Rather than being built on top of specific components and tables, application core logic and business rules are saved as metadata into the Force.com store.

Both application structure and application data are stored in the store.



A runtime engine executes application logic by retrieving its metadata and then performing the operations on the data.

Although running in isolated containers, different applications logically share the same database structure, and the runtime engine executes all of them uniformly.

A full-text search engine supports the runtime engine.

This allows application users to have an effective user experience despite the large amounts of data that need to be crawled.

The search engine maintains its indexing data in a separate store and is constantly updated by background processes triggered by user interaction.

APEX

Users can customize their application by leveraging the “native” Force.com application framework or by using programmatic APIs in the most popular programming languages.

The application framework allows users to visually define either the data or the core structure of a Force.com application, while the programmatic APIs provide them with a more conventional way for developing applications that relies on Web services to interact with the platform.

Customization of application processes and logic can also be implemented by developing scripts in APEX.

This is a Java-like language that provides object-oriented and procedural capabilities for defining either scripts executed on demand or triggers.

APEX also offers the capability of expressing searches and queries to have complete access to the data managed by the Force.com platform.

Microsoft Dynamics CRM

The solution implemented by Microsoft for customer relationship management Available either for installation on the enterprise's premises or as an online solution priced as a monthly per-user subscription.

The system is completely hosted in Microsoft's datacenters across the world and offers to customers a 99.9% SLA, with bonus credits if the system does not fulfill the agreement.

Each CRM instance is deployed on a separate database, and the application provides users with facilities for marketing, sales, and advanced customer relationship management.

Online features can be accessed either through a Web browser interface or programmatically by means of SOAP and RESTful Web services.

This allows Dynamics CRM to be easily integrated with both other Microsoft products and line-of-business applications.

Can be extended by developing plug-ins that allow implementing specific behaviors triggered on the occurrence of given events.

Can also leverage the capability of Windows Azure for the development and integration of new features.



Social networking

Social networking applications have grown considerably in the last few years to become the most active sites on the Web.

To sustain their traffic and serve millions of users seamlessly, services such as Twitter and Facebook have leveraged cloud computing technologies.

The possibility of continuously adding capacity while systems are running is the most attractive feature for social networks, which constantly increase their user base.

Points covered

Part III : Cloud Computing Applications

- ✓ Healthcare: ECG Analysis in the Cloud
- ✓ Biology: Protein Structure Prediction
- ✓ Geosciences: Satellite Image Processing
- ✓ Business and Consumer Applications:
 - ✓ CRM and ERP
 - ✓ Social Networking

Google Cloud Application: Google App Engine.

Overview of OpenStack architecture.



Google Cloud Applications

Google Applications Utilizing Cloud

Google App Engine

Google Applications Utilizing Cloud

Google's domain is constructed on the World Wide Web advertising.

In 2010, 96% of its \$29 billion income came from online ads.

Google deals subscriptions to enterprises, applying its web experience to market conventionally controlled by a very distinct kind of Software Company.



Like Google's search engine, Google Apps is constructed on top of a highly distributed infrastructure that separates both data and code over many servers and data centres.

This consistent back-end is conceived in order that if one data centre goes down, another can directly step into the breach.

Google values custom-built devices that permit it to improve services without taking them.

Google Applications Utilizing Cloud : Gmail

Gmail makes organizing the e-mail system so straightforward and efficient.

Gmail boasts 25 GB of storage per client, mighty spam filtering, BlackBerry and Outlook interoperability and a 99.9% uptime SLA (Service Level Agreement).

Gmail : Why / How?

- **E-mail, IM, voice and video chat:** Each client gets 25 GB of e-mail and IM storage.
- **Anytime, any location, get access to your e-mail:** Gmail is securely powered by the WWW, so you can be creative from your table, on the street, at home and on your wireless telephone, even when you are offline.
- **Sync with Android, iPhone and BlackBerry:** Get the advantages of Apps on premier wireless platforms.
- **Search and find e-mails instantly:** Spend short time in managing e-mail and locate e-mails quickly with Google-powered search to your inbox.
- **Get less spam:** Gmail spam is powerful filtering which assists you to concentrate on important ones.

Google Calendar

With Google Calendar, it is so simple to organize tasks all in one place.

- **Easily schedule appointments:** Join multiple calendars to preview when people are available. Google Calendar drives requests and organizes them.
- **Integrate with e-mail system:** Google Calendar is incorporated into Gmail and interoperable with popular calendar applications.
- **Share task calendars:** Calendars can be distributed company-wide or with chosen coworkers. A variety of distributing consent controls assist to sustain security and privacy.
- **Access with your wireless device:** View and edit happening minutia, add new happenings and ask for visitors on wireless devices like the BlackBerry and iPhone. Even obtain calendar notifications by SMS.
- **Publish calendars:** Publicize external business happenings by announcing a calendar to make it searchable in the Google Calendar gallery. Easily embed calendars into web pages.

Google Docs

An easy-to-use online word processor, spreadsheet and production reviewer that endows to perceive, store and share instantly and securely and cooperate online in less time.

Users can conceive new articles from the rub or upload living articles, spreadsheets and presentations.

There are no software programs to download and all your work is retained securely online and can be accessed from any computer.

Google Docs

Works over other operating systems: Google Docs works in the browser on PC, Mac and Linux computers and carries well-liked formats, for example, .doc, .xls, .ppt and .pdf.

Easily upload and share files: Files retained on Google Docs are habitually accessible and backed-up online.

Secure access to controls: Administrators can organize document distributing permissions system-wide and article proprietors can share and revoke document access at any time.

Google Sites

Google Sites are the easiest way to make knowledge accessible to the population who want speedy, up-to-date access.

People can work concurrently on a site to add record supplements, knowledge from other Google requests for paid jobs (like Google Docs, Google Calendar, YouTube and Picasa) and new free-form content.

Google Sites are accessible from any internet-bound computer.

Google Sites (2)

- **Organize knowledge in a central place:** Use Google Sites to centralize written material, spreadsheets, demonstrations, videos, slideshows and more to aid to retain the teams organized.
- **Anytime, everywhere access:** Google Sites are securely propelled by the web, so you can gather pages from your office desk, on the move, at home and on your portable phone.
- **Works through various operating systems:** Google Sites work on the browser installed on PC with Mac and Linux computers.
- **System and site-level security controls:** Administrators can supervise site sharing permissions through the enterprise and authors can share and revoke file access at anytime.

Google Groups

The Google Groups service creates a Google Group which is a user-owned group.

Google Groups not only sanctions us to supervise and record the mailing list, but in addition gives a manner for accurate communication and collaboration with gathered members.

Unlike other free mailing list services, Google Groups bids lavish storage fixes, customizable pages and unique organization options.

Google Groups are all about aiding users to bind with peoples, to access knowledge and convey effectually over computer communication and on the web.

- **Fast setup:** Create and supervise gatherings without burdening IT.
- **Sharing with a group:** Employees can share docs, calendars, sites, divided folders and videos with a gathering instead of individuals.
- **Searchable archives:** Group members can access and explore archives of posted items, conveyed to their e-mail lists to expeditiously find topics of interest.

Google Video

The Google Video index is the most comprehensive on the WWW, comprising millions of videos indexed and obtainable for viewing.

Using Google Video, one can explore and watch an ever-growing accumulation of video presentations, cinema clips, videos tunes, documentaries, private productions and more from all over the web.

- **A video channel for your business:** Video sharing makes valued communications like inside training and company announcements more engaging and effective.
- **Keep videos secured and private:** Employees can securely share videos with co-workers without uncovering private information.
- **Anytime, everywhere access:** Google Video is securely propelled by the web, so you can view videos from your office desk, on the road and at home.
- **Works in various operating systems:** Google Video works on a browser installed on PC with Mac and Linux OS.

Points covered

Part I : Amazon Web Services (AWS)

- ✓ Amazon Web Services and Components
- ✓ Amazon Simple DB
- ✓ Elastic Cloud Computing (EC2)
- ✓ Amazon Storage System
- ✓ Amazon Database services (Dynamo DB)

Points covered

Part II : Microsoft Cloud Services

- ✓ Azure core concepts
- ✓ SQL Azure
- ✓ Windows Azure Platform Appliance

Points to cover

Part III : Cloud Computing Applications

- ✓ Healthcare: ECG Analysis in the Cloud
- ✓ Biology: Protein Structure Prediction
- ✓ Geosciences: Satellite Image Processing
- ✓ Business and Consumer Applications:
 - ✓ CRM and ERP
 - ✓ Social Networking

Google Cloud Application: Google App Engine.

Overview of OpenStack architecture.

Google App Engine

Google's stimulating application development and hosting platform in the cloud

With it, the client can construct and establish web applications on Google's scalable high-traffic infrastructure.

App Engine carries apps written in Python or Java and they will execute on servers that use the identical expertise that forces Google's websites for pace and reliability.



App Engine applications are simple to construct and scale as the traffic and data grows.

To maintain App Engine, there are no servers available.

It helps the user to upload the application.

Google App Engine devotes you to get access to the identical construction blocks that Google values for its own applications.

It makes it simpler to construct an application that runs reliably, even under a hefty load and with a large amount of data.

- 
- The development environment contains the following features:
- Dynamic WWW assisting with full support of extensive web technologies.
 - Persistent storage with queries, sorting and transactions that are powered by Bigtable and GFS.
 - Scalability and load balancing are done automatically.
 - Google APIs for authenticating users and dispatching e-mail.
 - Fully boasted localized development environment.

Google App Engine bundles the construction blocks and takes care of the infrastructure stack, departing you more time to aim on composing code and advancing your application.

Google Apps for Business

Powerful cloud-based messaging and collaboration tools are offered by Google Apps for Business, from tiny organizations to huge organizations.

Google Apps is 100% hosted by Google, which decreases IT charges, minimizes up-keeping and management and simplifies primary setup.

With Google Apps for Business, client gets:

- Customized e-mail addresses
- Mobile e-mail, calendar and IM access
- No added hardware or programs
- Industry-leading spam filtering
- 24/7 internet note and telephone support
- 99.9% uptime assurance
- Dramatic cost savings

Choosing Google Apps not only saves money, but also saves an unbelievable amount of time. If the entire IT group focuses on methods and forward-thinking that can really advance the way enterprise operates.

Google Apps for Education

Google Apps for Education boasts worth that is yet to be agreed in the world of cloud-based messaging and collaboration.

For \$0/user/year, school, employees and scholars of informative organizations at all levels can leverage this huge set of customizable connection and collaboration tools.

Tools like Google Sites and Google Groups are ready-made for the world of learning, endowing data exchange and direction at solely new levels.

Google Apps adds the newest technologies and establishes best practices for data-centre administration, network application security and data integrity.

Google Apps for Education (1)

Eight ways how Google Apps advantages the campus are listed herewith:

- Students will love you for it
- Free up your IT
- Easy to deploy
- Save money
- Google defends your privacy
- Security as powerful as Google
- Innovation in real-time
- Collaborate globally



Google Apps for Government

Google Apps for Government presents all of the identical advantages that Google Apps for Business does, but with a supplemented level of security that stands up to even the largest levels of government standards.

With Google Apps, the government department benefitted from the scale and redundancy of distributed data centres around the globe.

Points covered

Part III : Cloud Computing Applications

- ✓ Healthcare: ECG Analysis in the Cloud
- ✓ Biology: Protein Structure Prediction
- ✓ Geosciences: Satellite Image Processing
- ✓ Business and Consumer Applications:
 - ✓ CRM and ERP
 - ✓ Social Networking
- ✓ Google Cloud Application: Google App Engine.
Overview of OpenStack architecture.

Overview of OpenStack architecture

What is OpenStack?

OpenStack is a cloud operating system that controls large pools of compute, storage, and networking resources throughout a datacenter, all managed through a dashboard that gives administrators control while empowering their users to provision resources through a web interface.

https://docs.openstack.org/2025.1/?_ga=2.169208417.1146438738.1743725397-255524807.1743725397

Open Source Cloud Platforms and Industry Standards

There are several open source cloud management platforms forming and continuously improving in the industry:

OpenStack and CloudStack

two most significant in the industry

Some of these cloud management platforms have broad industry support but might lag behind in terms of the features and functionality provided by commercial cloud management platforms.

Organizations that want to evaluate or deploy this class of cloud management system should consider the pros and cons of using open source versus commercially supported software platforms.

Open source's largest benefit is avoiding software vendor lock-in

OpenStack

An open source project with the largest community, code contributors, and cloud provider and system integrator involvement

Main goal : to support interoperability between cloud services while enabling enterprises to create Amazon-like cloud services.

A combination of modules that you can use to build, host, and operate your own cloud.

OpenStack modules are available to provide IaaS VMs, object and block storage, networking, identity, and many other services.

There is an OpenStack module for cloud management automation and orchestration called *Heat*

Another module called *Horizon* that is a customer-facing self-service configuration portal.



Being an open source project, all developed source code is submitted to the OpenStack committee, with the combined code being released to the public free of charge.

Notable founders and adopters of OpenStack include RackSpace, NASA, Hewlett-Packard, and IBM—each adding to and customizing OpenStack for their customer deployments and integration into other cloud management platforms and providers.

Beyond a cloud management platform, OpenStack is also seen as an industry standard for application interfaces, interoperability between cloud providers, and eventually will at least influence software-defined networking and datacenters in the future.

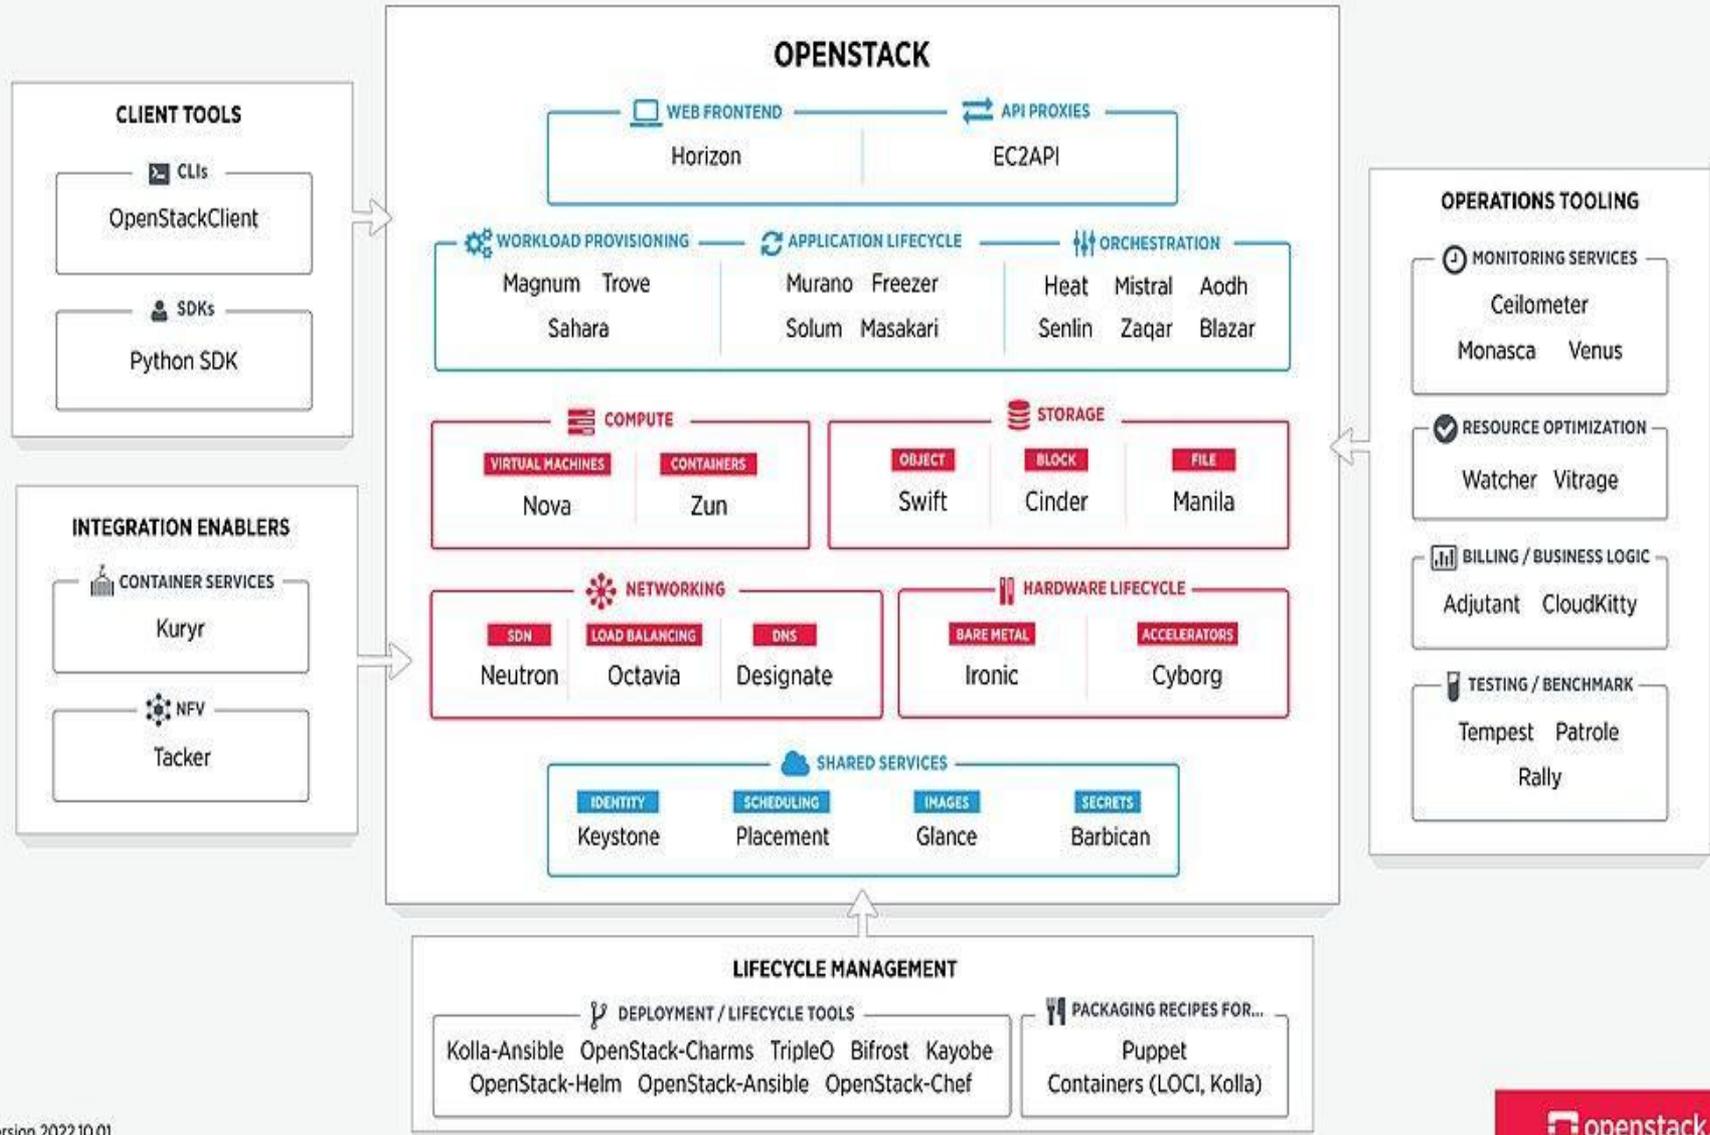
Given the quantity and well-known industry companies that have committed to OpenStack, it is expected to dominate as the industry open source for cloud and API integration between clouds.

OpenStack

A cloud operating system that controls large pools of compute, storage, and networking resources throughout a datacentre, all managed and provisioned through APIs with common authentication mechanisms.

Provides an Infrastructure-as-a-Service (IaaS) solution through a variety of complementary services.

Each service offers an Application Programming Interface (API) that facilitates this integration.





<https://www.openstack.org/software/project-navigator/openstack-components#openstack-services>

<https://docs.openstack.org/contributors/common/introduction.html>

IaaS Implementation Using OpenStack

IaaS can be implemented by integrating major components of OpenStack which are as follows:

- 1. Horizon:** Horizon is a web-based interface for managing OpenStack services. It provides a **graphical user interface** for operations such as launching instances, managing networking and setting access controls.
- 2. Keystone:** Keystone is the centralized **identity** service that provides authentication and authorization for other services. Keystone also provides a central catalog of services running in a particular OpenStack cloud.

IaaS Implementation Using OpenStack (1)

3. **Networking (neutron):** OpenStack Networking provides **connectivity** between the interfaces of other OpenStack services, such as Nova. OpenStack Networking is a pluggable architecture, users can create their own networks, control traffic, and connect servers to other networks.

4. **Cinder:** Cinder is a service that manages storage volumes for virtual machines. This is persistent **block storage** for the instances running in Nova.

IaaS Implementation Using OpenStack (2)

5. **Nova: Compute** nodes form the resource core of the OpenStack Compute cloud, providing the processing, memory, network and storage resources to run instances. Nova is a distributed component and interacts with Keystone for authentication, Glance for images and Horizon for web interface.

6. **Glance:** Glance service that acts as a registry for **virtual machine images** to allowing users to copy server images for immediate storage. Images can be used as templates when setting up new servers.



7. **Swift:** Swift service providing **object storage** which allows users to store and retrieve files. Swift architecture is distributed to allow for horizontal scaling, and to provide redundancy as failure proofing

8. **Manila :** The OpenStack **Shared File Systems** service (manila) provides file storage to a virtual machine. The Shared File Systems service provides an infrastructure for managing and provisioning of file shares. The service also enables management of share types as well as share snapshots if a driver supports them.



Trove: d/b management

Sahara: The **Data processing service** for OpenStack (sahara) aims to provide users with a simple means to provision data processing (Hadoop, Spark) clusters by specifying several parameters like Hadoop version, cluster topology, node hardware details and a few more. After a user fills in all the parameters, the Data processing service deploys the cluster in a few minutes. Sahara also provides a means to scale already provisioned clusters by adding or removing worker nodes on demand.

Course Outcome

CO4: Develop and deploy applications on Cloud