



MASARYKOVA UNIVERZITA

PV213 Enterprise Information Systems in Practice

07 – Architecture of the EIS in the cloud



MASARYKOVA UNIVERZITA

Tento projekt je spolufinancován Evropským sociálním fondem a státním rozpočtem České republiky.



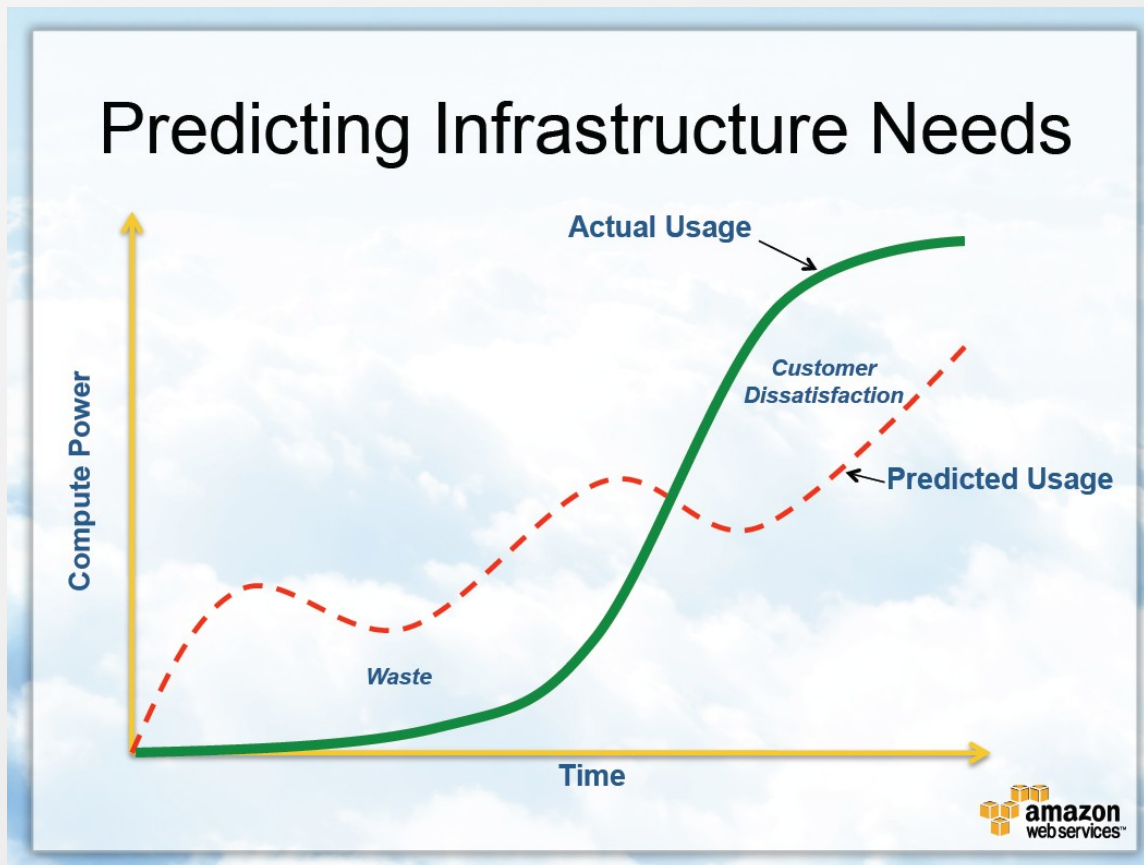
INVESTICE DO ROZVOJE VZDĚLÁVÁNÍ

Tento projekt je spolufinancován Evropským sociálním fondem a státním rozpočtem České republiky.



INVESTICE DO ROZVOJE VZDĚLÁVÁNÍ

Cloud Computing motivation



How can my application grow with customers?

How can I pay only for what I use?

What is Cloud Computing?

Cloud computing is Internet-based computing, whereby shared resources, software and information are provided to computers and other devices on-demand, like the electricity grid.

Five Attributes of Cloud Computing

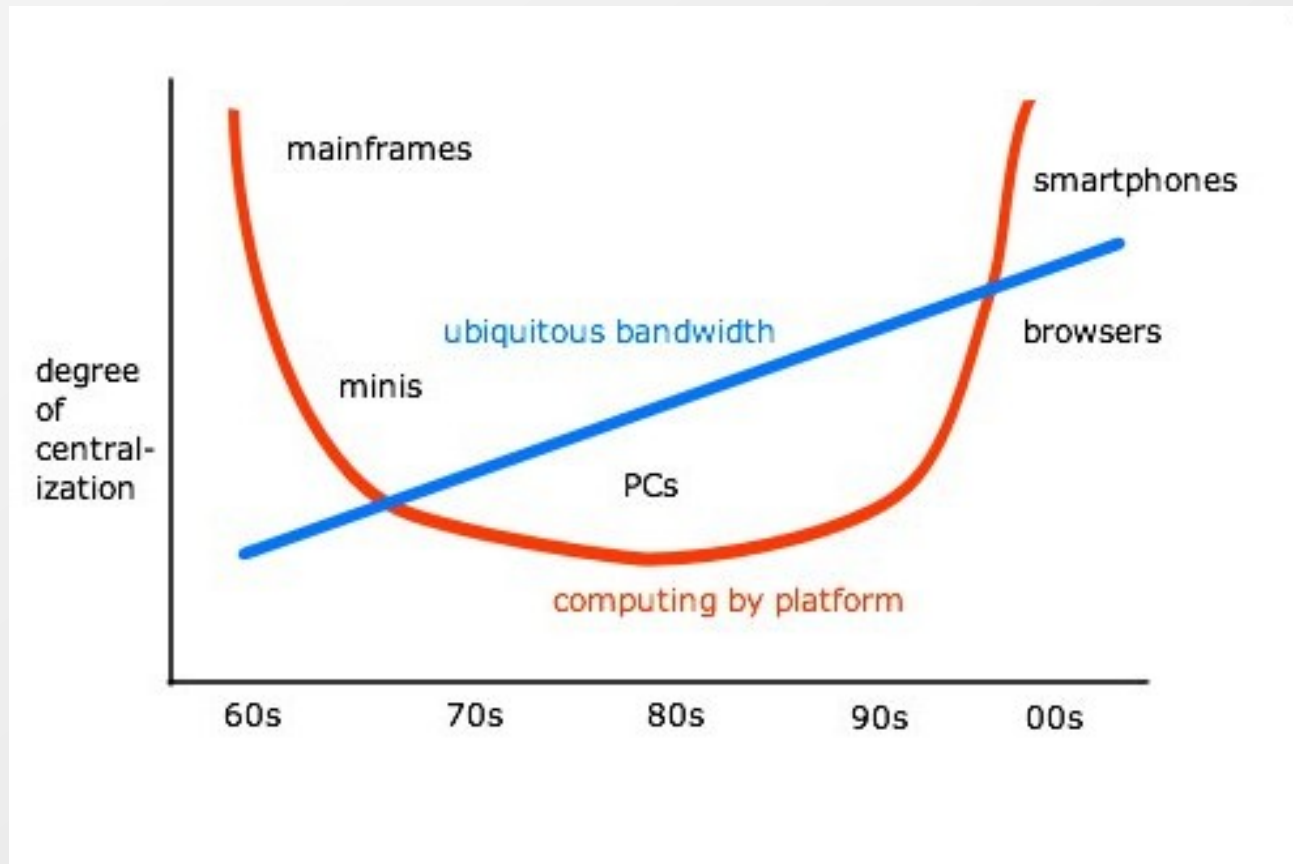
- ❑ Service based
- ❑ Scalable and elastic
- ❑ Shared
- ❑ Metered by use
- ❑ Uses internet technologies



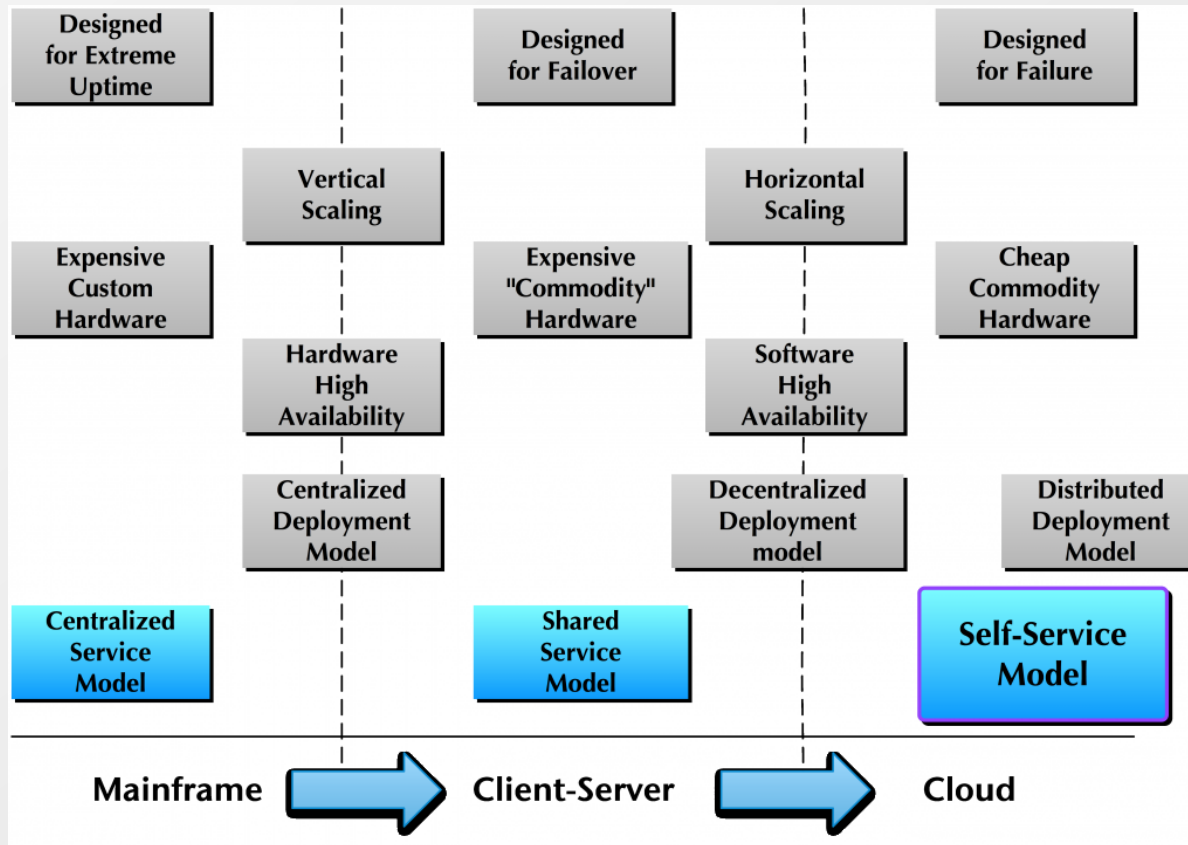
History

1960s	First ideas, book The Challenge of the Computer Utility	2008	Google App Engine Microsoft Azure Platform
1998	Akamai	2009	Google Apps VMware vCloud Express
1999	Salesforce.com	2010	Microsoft Office WebApps
2002	Google Search API Amazon Web Services	2011	IBM SmartCloud Microsoft Office 365
2004	Facebook	2012	Google Compute Engine Oracle Cloud HP Converged Cloud
2006	Amazon Elastic Compute Cloud Rackspace Cloud Twitter	2013	SAP HANA Enterprise Cloud Microsoft Azure Pack
2007	Dropbox Heroku		


Centralization / decentralization in history



Mainframe versus client-server versus cloud



Cloud Services Stack I



Software
as a Service

Platform
as a Service


Infrastructure
as a Service

Cloud Services Stack II

- **Virtualization** is a basis for all cloud services
 - Everything runs on virtualized HW
- **Infrastructure** defines basic available services
 - Operating system
 - Simple services
- **Platform** allows to build applications fast
 - Compatibility!
- **Application** is an only part visible for end users
 - Important is the value for end user


Types of cloud platforms

- ❏ **Public cloud**
 - ❏ “Classical” model
 - ❏ The same functionality for all
- ❏ **Private cloud**
 - ❏ Used only by one organization
- ❏ **Community cloud**
 - ❏ Functionality is used by several parties which share the same interest (e.g. area of business)
- ❏ **Hybrid cloud**
 - ❏ Use of two or more clouds (public, private, community) together as a one platform



Amazon Web Services (AWS)


Amazon is not only a web shop



Amazon is a leader in IaaS and PaaS
AWS introduced in 2002

- ❑ Online services for other web sites
- ❑ Accessing web shop data
- ❑ Evolved into big amount of services from different areas
- ❑ Available worldwide
- ❑ Billed on usage
- ❑ Accessed via HTTP, REST and SOAP
- ❑ Amount of services increase every year


Amazon Web Services against other cloud players



Amazon Web Services - Regions and availability zones I

Available regions (March 2015)

- ❏ US East - Northern Virginia (March 2006)
- ❏ US West - Northern California (December 2009)
- ❏ US West - Oregon (November 2011)
- ❏ EU - Ireland (November 2007)
- ❏ EU - Frankfurt (October 2014)
- ❏ Asia-Pacific - Singapore (April 2010)
- ❏ Asia-Pacific - Tokyo (April 2011)
- ❏ Asia-Pacific - Sydney (November 2012)
- ❏ China - Beijing (coming soon, in beta)
- ❏ South America - Sao Paulo (December 2011)
- ❏ AWS GovCloud - Special region for US government in Oregon (August 2011)



Amazon Web Services - Regions and availability zones II

Availability zones


- ❑ Distinct locations within one region
- ❑ Availability zones in the region are physically separated
- ❑ Even disasters like flooding should affect only single availability zone
- ❑ Transfers between availability zones in the same region are cheap

AWS Edge Locations

- ❑ Content delivery network for distributing static content
- ❑ Distributed around the all world (March 2015)
 - ❑ North America - 14 locations
 - ❑ South America - 2 locations
 - ❑ Europe - 10 locations
 - ❑ Asia - Pacific - 11 locations

Amazon Web Services - Regions and availability zones III

- ❑ Each region has two or more availability zones
- ❑ You put your application which is “nearest” to you
- ❑ You can use more region but you have to pay for usage of each of them
- ❑ Not all services are available in all regions



Amazon Web Services - Portfolio I

Compute

- ❑ Elastic Compute Cloud (EC2, August 2006)
 - ❑ Virtualized OS in Amazon's infrastructure
- ❑ Auto Scaling (May 2009)
 - ❑ Allows to automatically scale EC2 capacity up or down
- ❑ Elastic Load Balancing (May 2009)
 - ❑ Distributes requests to multiple EC2 instances
- ❑ EC2 Container Service
 - ❑ Docker container management service (November 2014)
- ❑ AWS Lambda
 - ❑ Runs code in response to events, pays only execution (November 2014)



Amazon Web Services - Portfolio II

Storage and Content Delivery

- ❏ Simple Storage Service (S3, March 2006)
 - ❏ Key-value storage for big objects (images, video, documents, etc.)
- ❏ Glacier (August 2012)
 - ❏ Low cost storage service for data archiving and backup
- ❏ Elastic Block Store (EBS, August 2008)
 - ❏ Block device mounted by EC2 instances (virtual hard drive)
- ❏ Import/Export (May 2009)
 - ❏ Allows to import/export huge amount of data from/to portable device
- ❏ Storage Gateway (January 2012)
 - ❏ Integrates on-premises IT environments with Cloud storage

Amazon Web Services - Portfolio III

Databases

- Relation Database Service (October 2009)
 - Relational database in cloud (MySQL, Oracle, SQL Server, Postgre SQL)
- SimpleDB (December 2007)
 - Non-relational data store
- DynamoDB (January 2012)
 - NoSQL database
- ElastiCache (August 2011)
 - In-memory cache in the cloud

Amazon Web Services - Portfolio IV

Application Services I

- ❏ Simple Queue Service (SQS, July 2006)
 - ❏ Message queue in the cloud
- ❏ Flexible Payments Service (FPS, August 2007)
 - ❏ Payment service for developers
- ❏ Simple Notification Service (SNS, April 2010)
 - ❏ Sends notifications from the cloud
- ❏ Simple Email Service (SES, January 2011)
 - ❏ Scalable e-mail sending service

Amazon Web Services - Portfolio V

Application Services II

- ❏ Simple Workflow Service (SWF, February 2012)
 - ❏ Workflow service for coordinating application components
- ❏ CloudSearch (April 2012)
 - ❏ Managed search service
- ❏ Elastic Transcoder (January 2013)
 - ❏ Easy-to-use scalable media transcoding
- ❏ AppStream (November 2013)
 - ❏ Low-Latency Application Streaming

Amazon Web Services - Portfolio VI

Analytics

- ❏ Elastic MapReduce - EMR (April 2009)
 - ❏ Process vast amount of data (Hadoop framework)
- ❏ Redshift (November 2012)
 - ❏ Petabyte-scale data warehouse service
- ❏ Data pipeline (December 2012)
 - ❏ Orchestration Service for periodic, data-driven workflows
- ❏ Kinesis (November 2013)
 - ❏ Real-Time data stream processing

Amazon Web Services - Portfolio VII

Networking

- ❏ Virtual Private Cloud (VPC, August 2009)
 - ❏ Allows to create isolated services in virtual network
- ❏ Direct Connect (August 2011)
 - ❏ Connects your existing private data center with Amazon's cloud
- ❏ Route 53 (December 2010)
 - ❏ Scalable domain name system (DNS)

Content Delivery

- ❏ CloudFront (November 2008)
 - ❏ Content delivery service

Amazon Web Services - Portfolio VIII

Administration and Security I

- ❏ Management Console (January 2009)
 - ❏ Web-Based User Interface
- ❏ CloudWatch (May 2009)
 - ❏ Monitoring of cloud resources
- ❏ Identity and Access Management (IAM, September 2010)
 - ❏ Securely controls access to services and resources
- ❏ Trusted Advisor (Jan 2012)
 - ❏ Helps you provision your resources by following best practices

Amazon Web Services - Portfolio IX

Administration and Security II

- ❏ CloudHSM (March 2013)
 - ❏ Hardware-based Key Storage for Regulatory Compliance
- ❏ Key Management Service (November 2014)
 - ❏ Create and control encryption keys used to encrypt data
- ❏ CloudTrail (November 2013)
 - ❏ Logs API calls and delivers log files
- ❏ Directory Service (Oct 2014)
 - ❏ Active Directory
- ❏ Config (Nov 2014)
 - ❏ Detailed inventory of AWS resources with history changes

Amazon Web Services - Portfolio X

Deployment and Management

- ❏ Elastic Beanstalk (January 2011)
 - ❏ Deployment and management tool
- ❏ CloudFormation (February 2011)
 - ❏ Tool for creating templates of related AWS resources
- ❏ OpsWorks (February 2013)
 - ❏ Application management service
- ❏ CodeDeploy
 - ❏ Automates code deployments (November 2014)

Amazon Web Services - Portfolio XI

Mobile Services

- ❏ Congnito (July 2014)
 - ❏ Identity and data synchronization service
- ❏ Mobile Analytics (July 2014)
 - ❏ Collect, visualize and understand app usage

Enterprise Applications:

- ❏ WorkSpaces (November 2013)
 - ❏ Virtual Desktops in the Cloud
- ❏ WorkDocs (formerly Zocalo, July 2014)
 - ❏ Enterprise storage and sharing service

Amazon Web Services - Portfolio XII

Marketplace

- ❑ Marketplace (April 2012)
 - ❑ Find software that runs on EC2 and launch with 1-Click

Web Traffic

- ❑ Alexa Web Information Service (October 2004)
 - ❑ Web statistics
- ❑ Alexa Top Sites (January 2006)
 - ❑ Top internet site statistics

Workforce

- ❑ Mechanical Turk (November 2005)
 - ❑ Manage tasks made by human via cloud



Amazon Web Services - Compute Services I

- ❏ Elastic Compute Cloud (EC2) is infrastructure for running virtual machines
- ❏ EC2 is a core functionality for all customer specific computation
- ❏ EC2 works as IaaS
 - ❏ You can install everything what you want (licenses!)
- ❏ There exists pre-prepared AMIs (Amazon Machine Images) - virtual machines for “ready to use” (several thousand of AMIs exists in different categories)
- ❏ AMIs are available with following OSes:
 - ❏ Linux: Red Hat, SUSE, CentOS, Ubuntu, Fedora, FreeBSD, Debian, ...
 - ❏ Windows server: 2003 R2, 2008, 2008 R2, 2012

Amazon Web Services - Compute Services II

- ❑ AMIs also contain pre-prepared middleware or even applications
 - ❑ Databases, web servers, application development environments, application servers, video encoding & streaming, ...
 - ❑ SAP HANA, Oracle JD Edwards, Oracle PeopleSoft, ...
- ❑ Price of AMIs is AMI dependant
 - ❑ Some AMIs are free of charge
 - ❑ For other AMIs you pay per hours virtual machine (instance) is running

You can create your own EC2 instances with your required software!

Amazon Web Services - Compute Services III

- ❏ Amazon Elastic Compute Cloud represents core functionality for you app
- ❏ You deploy your application to given region
- ❏ For dynamic behavior use Auto Scaling functionality
 - ❏ You define rules for higher/lower workloads
 - ❏ Based on current statistics
 - ❏ Time scheduled
 - ❏ Good when you have peaks in usage
- ❏ Elastic Load Balancing allows to distribute requests to your app
 - ❏ Can increase availability of your application
 - ❏ Works together with Auto Scaling
- ❏ Elastic IP Addresses allows you to have virtual IP addresses
 - ❏ Public IP addresses mapped dynamically to given instance

Amazon Web Services - Compute Services IV

- ❏ You can monitor instances via CloudWatch
- ❏ Amazon guarantees 99.95% availability for region
- ❏ By default instances are not backed-up!
 - ❏ Don't store your data into the instance. Use e.g.
 - ❏ S3 for simple unstructured data
 - ❏ Database for structured data (keep scalability in mind)
 - ❏ Elastic Block Store (EBS - file system) - can be attached just to one EC2 instance
 - ❏ For storage you have to pay extra
- ❏ If you run several instances (e.g. in different regions) you pay for each
 - ❏ Prices in different regions are usually different
- ❏ You can order on-demand, reserved or spot instance

Amazon Web Services - Compute Services V

Amazon EC2 surroundings

- ❏ Around EC2 are management services
 - ❏ Auto scaling
 - ❏ Elastic load balancing
 - ❏ CloudWatch
 - ❏ Management console
 - ❏ ...



Amazon Web Services - Instance types

You have to choose hardware on which you want to run, e.g.

- ❑ Low cost t2.micro instance - 1 GiB memory, 1vCPU, EBS only storage
- ❑ General purpose m3.xlarge instance - 15 GiB memory, 4 vCPUs, 2x40 GiB SSD instance storage
- ❑ Compute optimized c3.2xlarge instance - 15 GiB memory, 8 vCPUs, 2x80 GiB SSD instance storage
- ❑ GPU g2.2xlarge instance - 15 GiB memory, 8 vCPUs, 1x60 SSD memory storage
- ❑ Memory optimized r3.8xlarge instance - 244 GiB memory, 32 vCPUs, 2x320 GiB SSD instance storage
- ❑ Storage optimized i2.4xlarge instance - 122 GiB memory, 16 vCPUs, 4x800 GiB SSD instance storage

Amazon Web Services - Simple Storage Service (S3) I

- ❑ Purpose of Simple Storage Service (S3) is to store static unstructured data (pictures, videos, documents, ...)
- ❑ Simple key-value storage with possible metadata organized into “buckets”
- ❑ Maximal size of each data (file) is 5 TB, unlimited number of objects
- ❑ Data accessible as <http://s3.amazonaws.com/bucket/key> (DNS CNAME supported - <http://yourdomain/bucket/key>)
- ❑ Data stored in regions, they never leave region if you don't require it
- ❑ Data can be versioned (different data with the same key) and encrypted
- ❑ Supports protocols HTTP and BitTorrent
- ❑ 99.999999999% durability, 99.99% availability
- ❑ With Reduced Redundancy Storage (RRS) durability and availability 99.99%

Amazon Web Services - Simple Storage Service (S3) II

- ❏ “Unexpected” behavior due to distributed nature
 - ❏ Immediate read after write can return “not available”
 - ❏ Immediate list of bucket content after write can return list without currently written object
 - ❏ Immediate read after re-write can still return “old” data
 - ❏ Immediate read after delete can still return “deleted” data
- ❏ Some web sites uses S3 for hosting static content
 - ❏ E.g. well known service DropBox is using Amazon S3 for storing data
- ❏ With Amazon Glacier you can reduce storage costs (with some restrictions)
 - ❏ Glacier is low-cost storage with long retrieval times (several hours)

Amazon Web Services - Database Services I

❏ Relational Database Service

- ❏ Allows to easily migrate your existing applications
- ❏ Support for MySQL, Oracle, Microsoft SQL Server or PostgreSQL
- ❏ By default you are allowed to run up to 40 instances (but you can ask for more)
- ❏ Size limit per instance is 3TB


❏ SimpleDB

- ❏ Non-relational data store, automatically indexes your data
- ❏ Each row can have different attributes
- ❏ Domain - item - attribute approach
- ❏ Limit is 10GB per domain, you can have 250 domains by default (but you can ask for more)

Amazon Web Services - Database Services II

SimpleDB (continues)

- 256 attributes per item, attribute length is limited to 1024 bytes
- Consistency
 - Eventually consistent reads (default)
 - Consistent reads
- Limited transactions (only conditional Put/Delete)
- Good for lower-scale workloads
- Older technology than DynamoDB



Amazon Web Services - Database Services III

❏ DynamoDB

- ❏ NoSQL (non-relational) database
- ❏ Fast predictable performance, stored all data on SSD
- ❏ No limits on amount of data, scales automatically
- ❏ By default you can order guaranteed throughput rate up to 10000 reads or writes per second (in exceptional cases you can order higher)
- ❏ Table - item - attribute approach
- ❏ Requires primary key for item
 - ❏ Primary key can be compound (hash and range type primary key)
- ❏ Allows to use secondary global indexes
- ❏ Item size limit is 400kB (sum of all attribute names and values)
- ❏ Not so flexible queries
- ❏ Supports batch processing

Amazon Web Services - Database Services IV


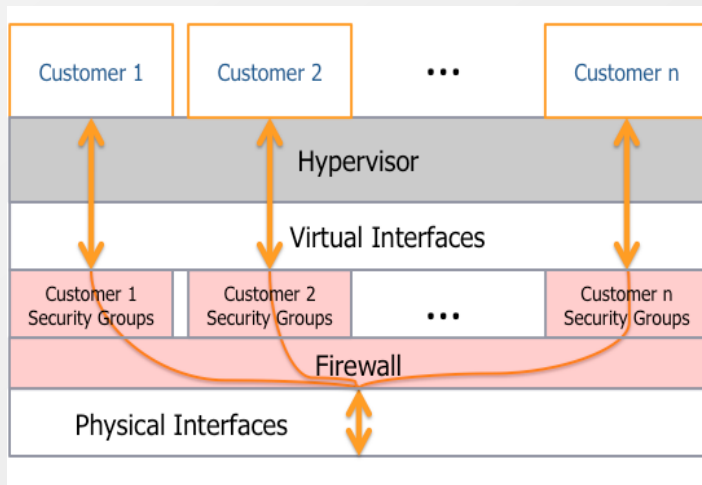
❏ DynamoDB (continues)

- ❏ Supports scalar data types: number, string, binary
 - ❏ Plus multi-valued types: number set, string set, binary set
- ❏ Consistency
 - ❏ Eventually consistent reads (default, better performance)
 - ❏ Consistent reads
- ❏ Integrates Elastic MapReduce
 - ❏ MapReduce is programming model for processing large data


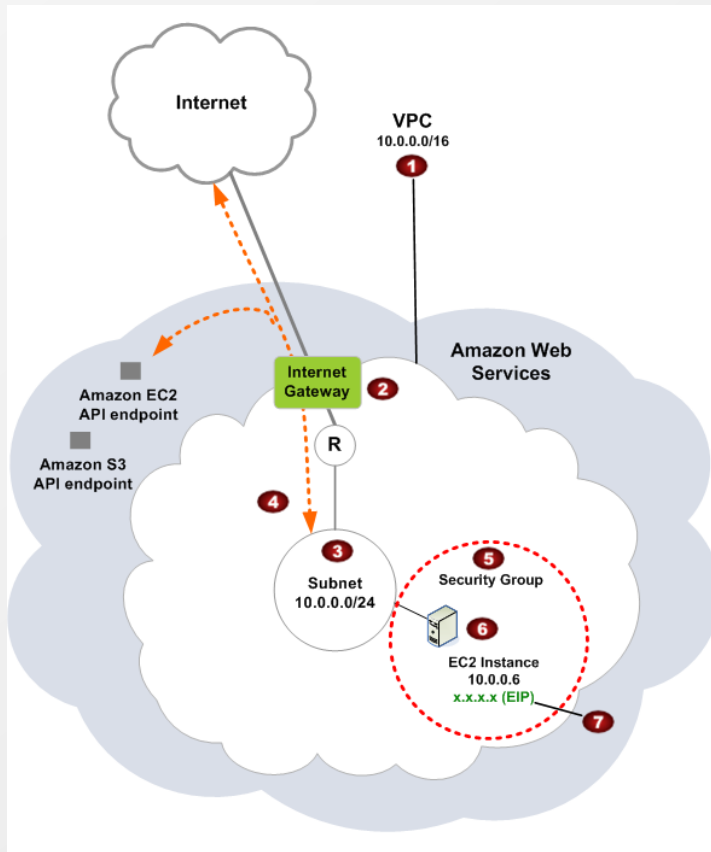
You can run also other databases like IBM DB2, Sybase, etc. in your own instance (they are not officially supported).

Amazon Web Services - Virtual Private Cloud

- Connection via VPN
- Isolated network access
- Subnets

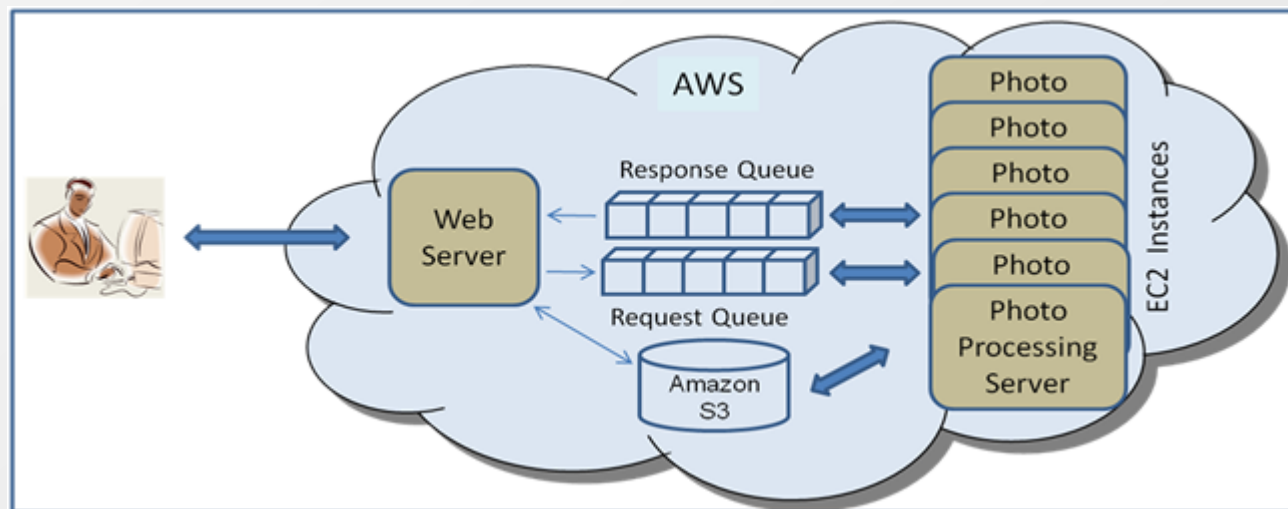


Amazon Web Services - Virtual Private Cloud, Security



Amazon Web Services - EC2 and S3 Example

Photo processing example



Amazon Web Services - Pricing I

Pricing

- ❏ Different for region, service, OS, ...
- ❏ Based on usage - you can have AWS account but you don't need to pay at all if you don't use Amazon's services
- ❏ Usually calculated per hour or per GB
- ❏ Some services are paid per item or per item and month (e.g. CloudWatch, Export/Import)
- ❏ Hours for EC2 instances means time when your instance is live!
- ❏ Prices are developing through years
- ❏ There are "action prices" all the time

Amazon Web Services - Pricing II

EC2 pricing depends on

- ❑ Operating system - Linux (Unix) / Windows (Linux is cheaper)
- ❑ Region
- ❑ Instance type (small / medium / large / ...)
- ❑ On-demand instance / reserved instance / spot instance
- ❑ Length of contract for reserved instances (1 year / 3 years)

S3 pricing depends on

- ❑ Region
- ❑ Amount of stored data
- ❑ Standard redundancy / reduced redundancy / glacier storage
- ❑ Number of GET, PUT, POST, ... requests
- ❑ Amount of data transferred

Amazon Web Services - Pricing III

New AWS customers get each month for one year (March 2015)

- ❏ 750 hours of EC2 running Linux/Unix or RHEL or SLES t2.micro instance usage
- ❏ 750 hours of EC2 running Microsoft Windows Server t2.micro instance usage
- ❏ 750 hours of Elastic Load Balancing plus 15 GB data processing
- ❏ 30 GB of Amazon Elastic Block Storage (EBS) plus 2 million IOs and 1 GB snapshot storage
- ❏ 5 GB of Simple Storage Service (S3), 20000 GET, 2000 PUT requests
- ❏ 15 GB of bandwidth out aggregated across all AWS services
- ❏ ...
- ❏ See <http://aws.amazon.com/free/> for details

Amazon Web Services - Pricing IV

Amazon pricing overview

☞ <http://aws.amazon.com/pricing/>

Amazon pricing for specific products

☞ <http://aws.amazon.com/ec2/pricing/>

☞ <http://aws.amazon.com/s3/pricing/>

☞ <http://aws.amazon.com/dynamodb/pricing/>

☞ ...

Simple Monthly Calculator and TCO calculator

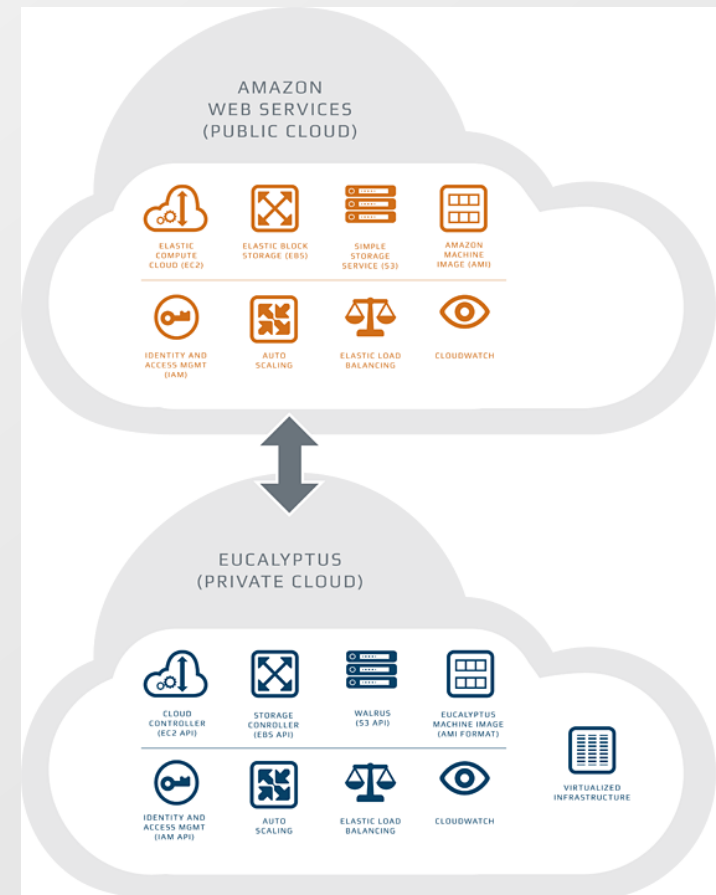
☞ <http://aws.amazon.com/calculator/>

☞ <http://aws.amazon.com/tco-calculator/>

Amazon Web Services - Alternatives

Open source alternative - Eucalyptus

- Runs some of AWSs under your Linux
- <http://www.eucalyptus.com>



Microsoft Azure Platform

- ❏ Microsoft's response to the cloud computing
 - ❏ Announced in October 2008
 - ❏ Commercially available in February 2010
 - ❏ IaaS, PaaS and partially SaaS
- ❏ Windows Azure OS as a “base OS”
 - ❏ Windows Server 2008 R2 for Guest OS 2.x
 - ❏ Windows Server 2012 for Guest OS 3.x
 - ❏ Windows Server 2012 R2 for Guest OS 4.x
- ❏ SLAs (monthly)
 - ❏ 99.99% for Traffic Manager (distribution of user traffic)
 - ❏ 99.95% for Cloud Services and Virtual Machines
 - ❏ 99.9% for storage, CDN, SQL Database, service bus, caching, ...



Windows Azure Data Centers (March 2015)

- ❑ 17 regions around the world
 - ❑ USA - Iowa, Virginia, Illinois, Texas, California
 - ❑ South America - Brazil
 - ❑ Europe - Netherlands, Ireland
 - ❑ Asia - Hong Kong, Singapore, Japan
 - ❑ Australia - Victoria, New South Wales

- ❑ 31 Content Delivery Network locations
 - ❑ USA - 10
 - ❑ Europe - 11
 - ❑ Asia - 8
 - ❑ Australia - 2



Windows Azure - Services overview I

Compute

- Virtual Machines
- Cloud Services
- Batch
- RemoteApp

Web and Mobile

- Web Sites
- Mobile Services
- Push Notifications
- Mobile Engagement

Networking

- Virtual Network
- ExpressRoute
- Traffic Manager

Data Services

- Storage (Table, BLOB)
- DocumentDB
- SQL Database
- Redis Cache
- StorSimple
- Azure Search

Windows Azure - Services overview II

Analytics

- HDInsight
- Machine Learning
- Stream Analytics
- Data Factory
- Event Hubs

Media and CDN

- Media Services
- CDN

Hybrid Integration

- BizTalk Services
- Service Bus
- Backup
- Site Recovery

Identity and Access Management

- Azure Active Directory
- Multi-Factor Authentication

Developer Services

- Visual Studio Online
- Application Insights

Windows Azure - Compute services

Virtual Machines

- IaaS - Run Windows or Linux, predefined or own image

Cloud Services

- Multi tier cloud applications
- One of the first services available in Microsoft Azure
 - Web role
 - Worker role

Batch

- Massive parallel batch processing

RemoteApp

- “Terminal Server” in the cloud

Windows Azure - Web and Mobile

☒ Web Sites

- ☒ Hosting web applications written in .NET, PHP, Java, Node.js or Python
- ☒ With WebJobs you can run also some background processing

☒ Mobile Services

- ☒ PaaS for building backend for mobile applications

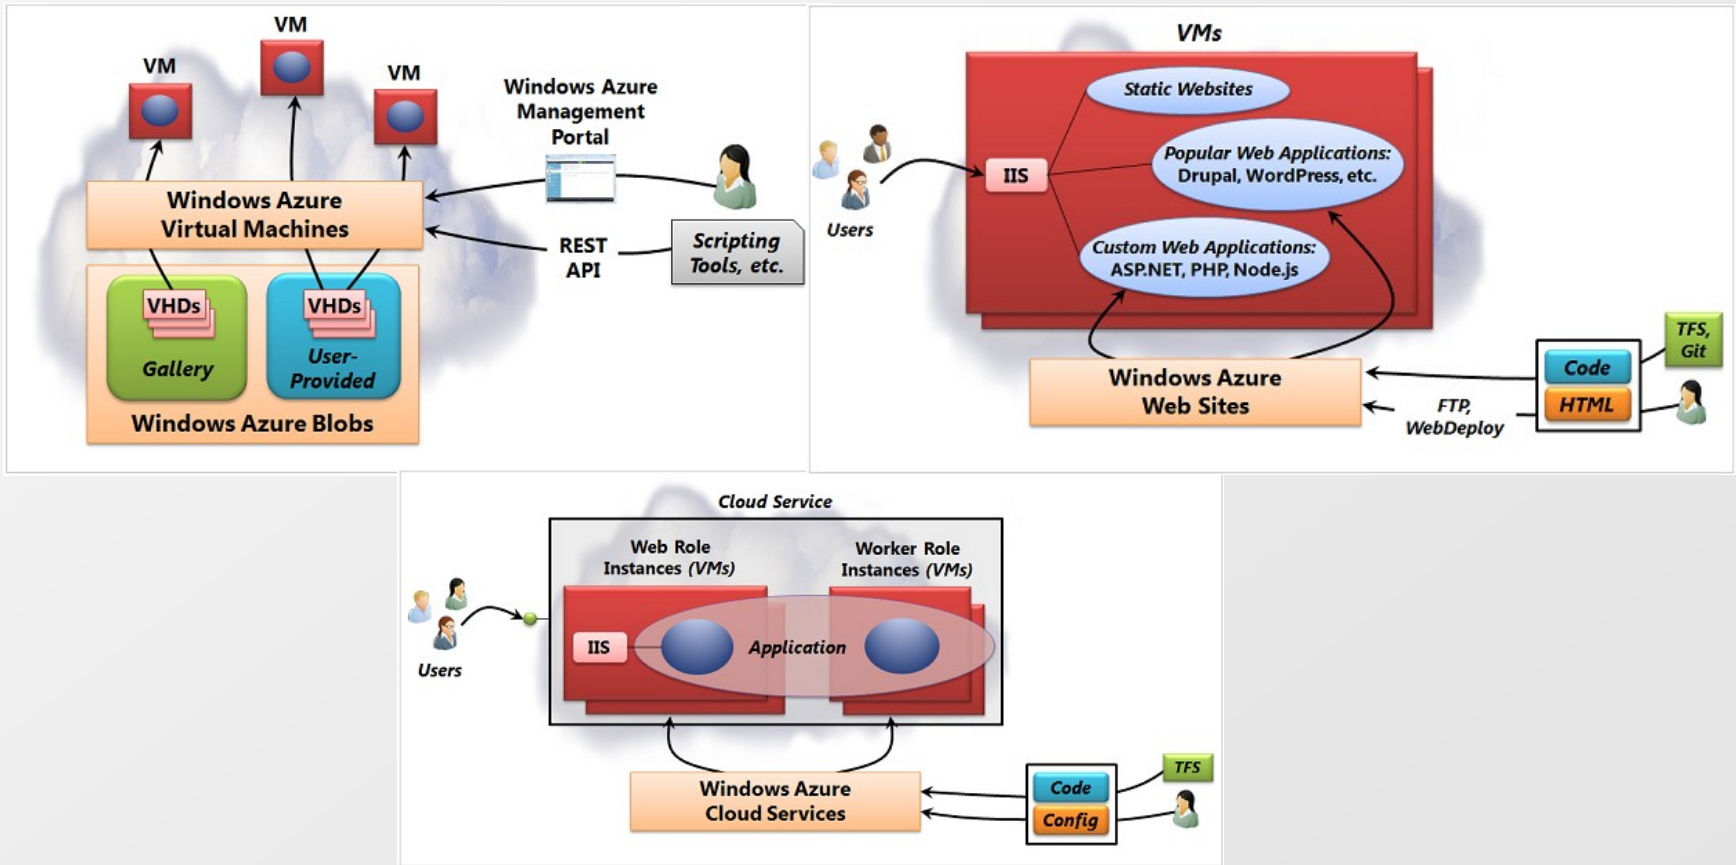
☒ Push Notifications

- ☒ Highly scalable push notification infrastructure to millions of devices

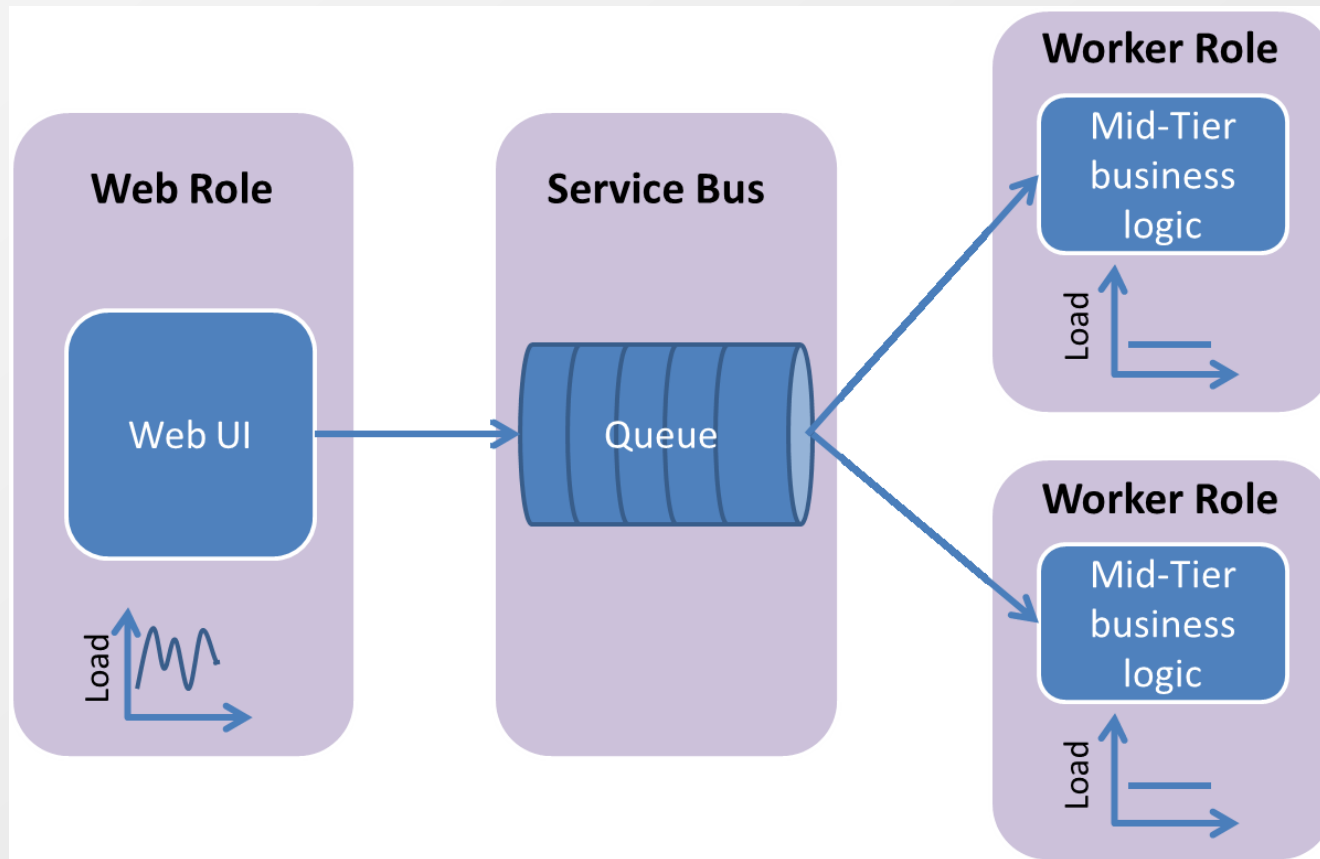
☒ Mobile Engagement

- ☒ Collect and analyze usage of the mobile application

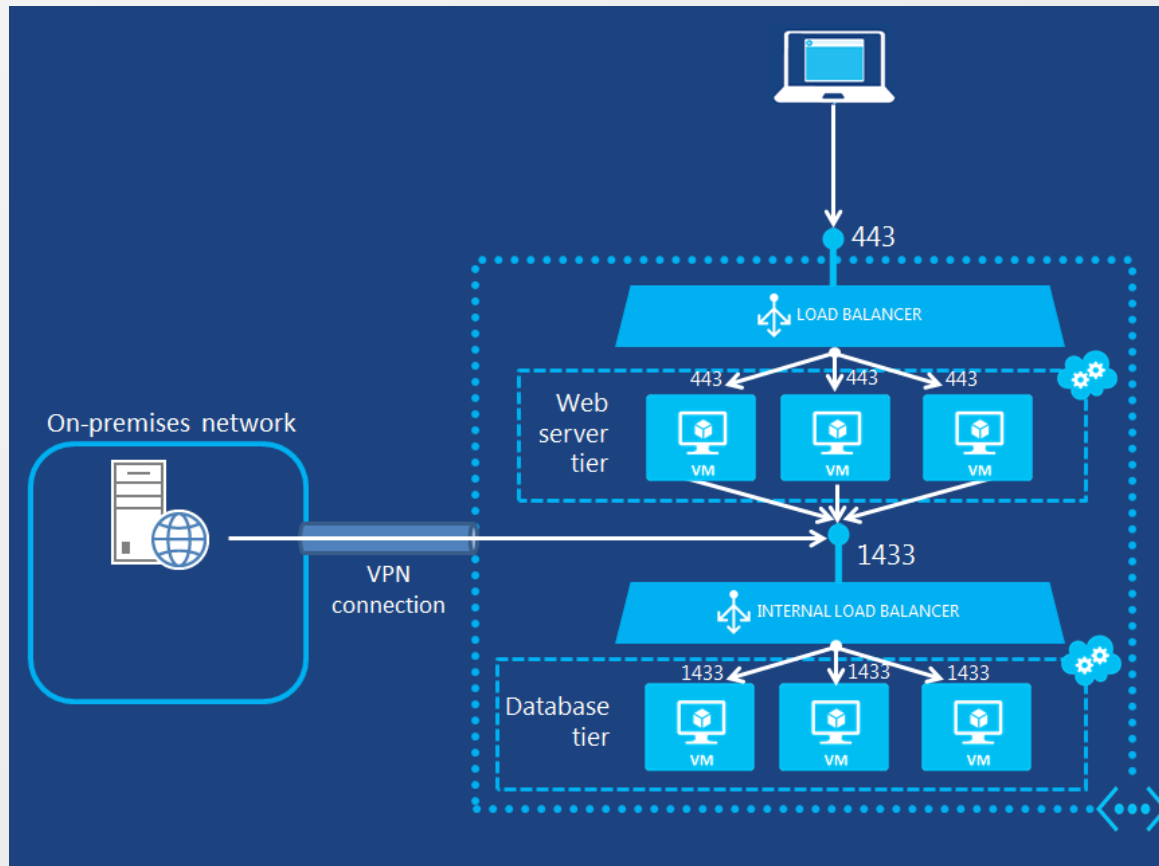
Windows Azure - VMs vs. Cloud Services vs. Web Sites



Windows Azure - Web role and worker role



Windows Azure - Load balancing explanation



Windows Azure - Data services I

- ❏ **BLOB (Binary Large Objects)**
 - ❏ Key-value storage for unstructured data (video, documents, etc.)
- ❏ **Table**
 - ❏ NoSQL key-attribute data store
- ❏ **Queue**
 - ❏ Reliable messaging for workflow processing and communication
- ❏ **File storage**
 - ❏ Storage for legacy applications using the standard SMB protocol
- ❏ **SQL Database**
 - ❏ MS SQL Server database in the cloud
 - ❏ Doesn't support all features of standard MS SQL Server

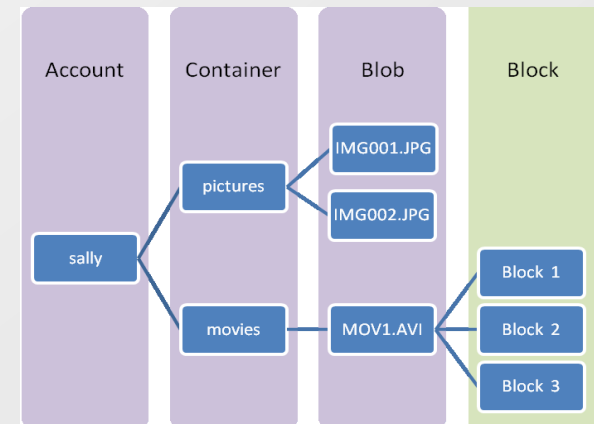
Windows Azure - Data services II

- ❏ **DocumentDB**
 - ❏ Document based NoSQL database
- ❏ **Redis Cache**
 - ❏ Open source distributed, scalable, in-memory cache
- ❏ **StorSimple**
 - ❏ Enterprise storage and data protection
- ❏ **Azure Search**
 - ❏ Search service

Microsoft Azure - Data Services - BLOB and Table

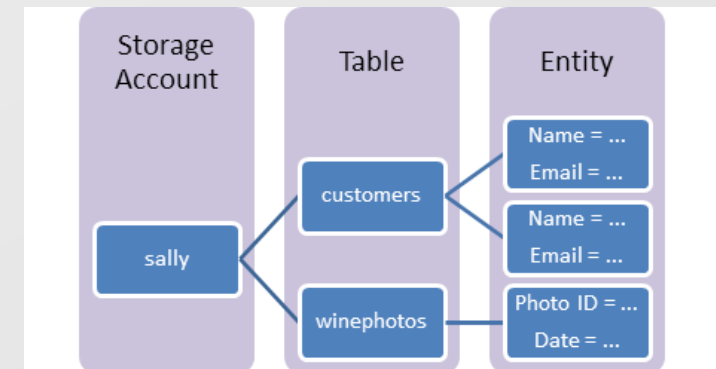
BLOB

- ❏ Maximal 500 TB of data per account
- ❏ Blob - Key-value storage with “block” support
 - ❏ Up to 200 GB for block blob
 - ❏ Up to 1 TB for page blob
 - ❏ Optimized for random read/write
 - ❏ <http://<storage account>.blob.core.windows.net/<container>/<blob>>



Table

- ❏ NoSQL indexed database
- ❏ Maximal entity (row) size is 1 MB
- ❏ Maximum 255 properties in entity (row)
- ❏ Maximal property size is 64 kB




Microsoft Azure - Data Services - Azure SQL Database

- ❏ Supports T-SQL
- ❏ Maximal 150 databases in each SQL Server
- ❏ Maximal limits per database
 - ❏ Web Edition - 5GB
 - ❏ Business Edition - 150GB
- ❏ For larger data you have to use sharding (split data into several databases)
- ❏ Every table must have clustered index
- ❏ Limits: no full text search, CLR types, replication, database mirroring, jobs, backup and restore, distributed transactions, ...

Microsoft Azure - Data partitioning

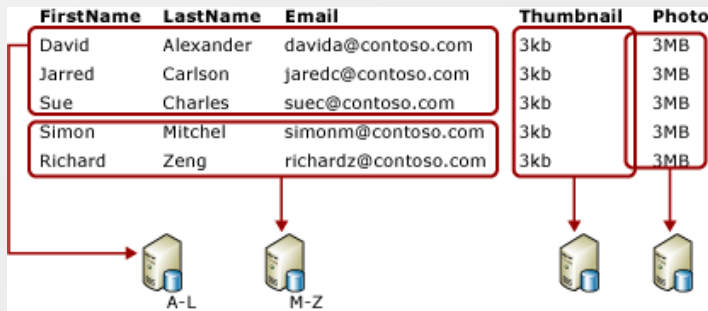
Horizontal partitioning



Vertical partitioning



Hybrid partitioning



Windows Azure - Networking services

❏ Windows Azure Virtual Network

- ❏ VPN between your intranet and Azure

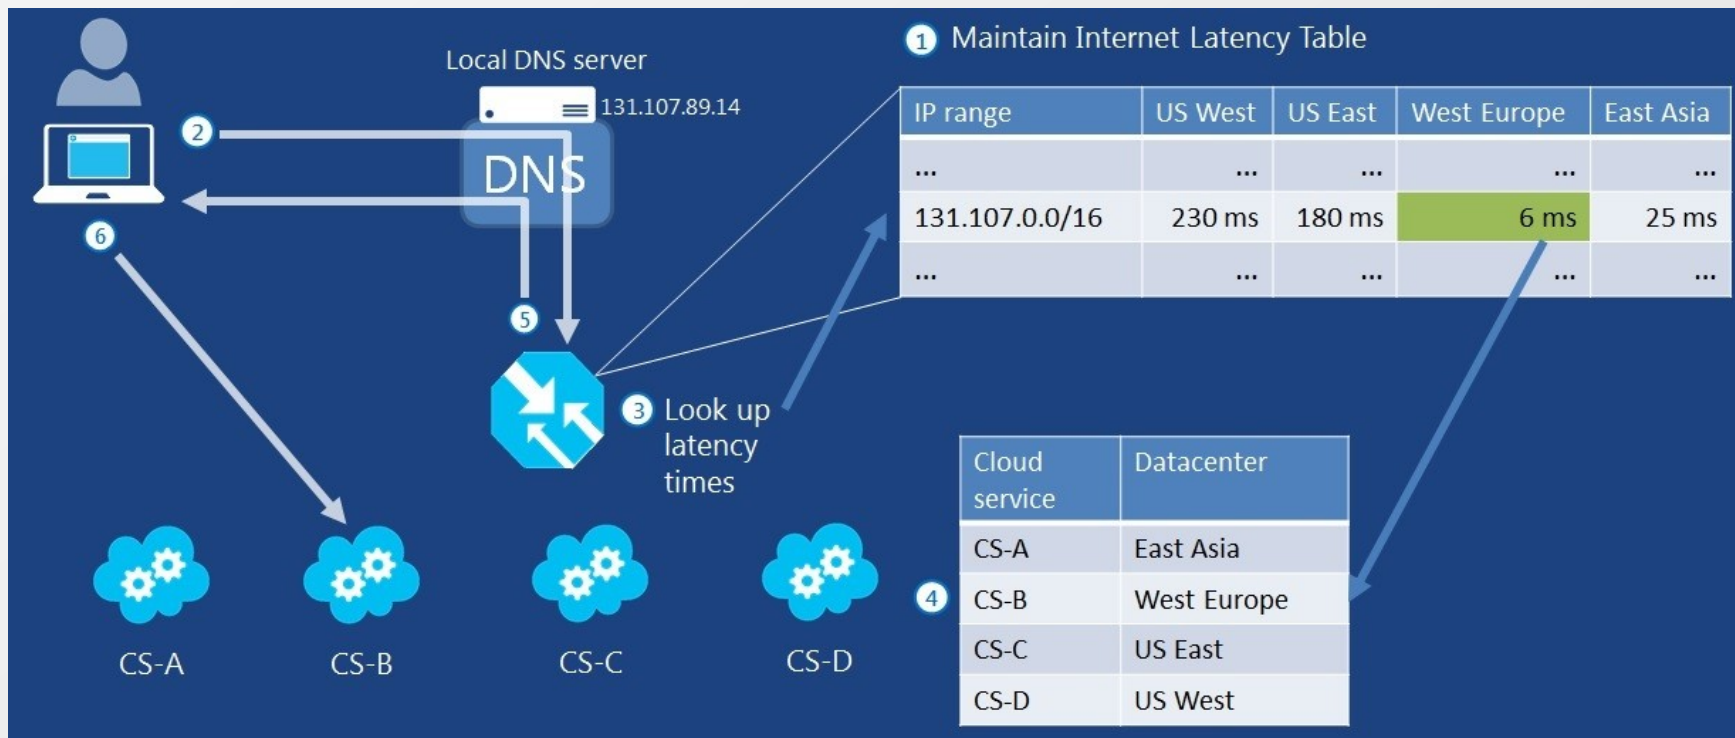
❏ Traffic Manager

- ❏ Controls distribution of user traffic to similar hosted services running in the same or different data center according to conditions like performance, price or compliance

❏ ExpressRoute



- ❏ Private (non-internet) connection to the Azure data centers
- ❏ Available only on limited number of locations with limited number of network providers

Windows Azure - Performance Load Balancing



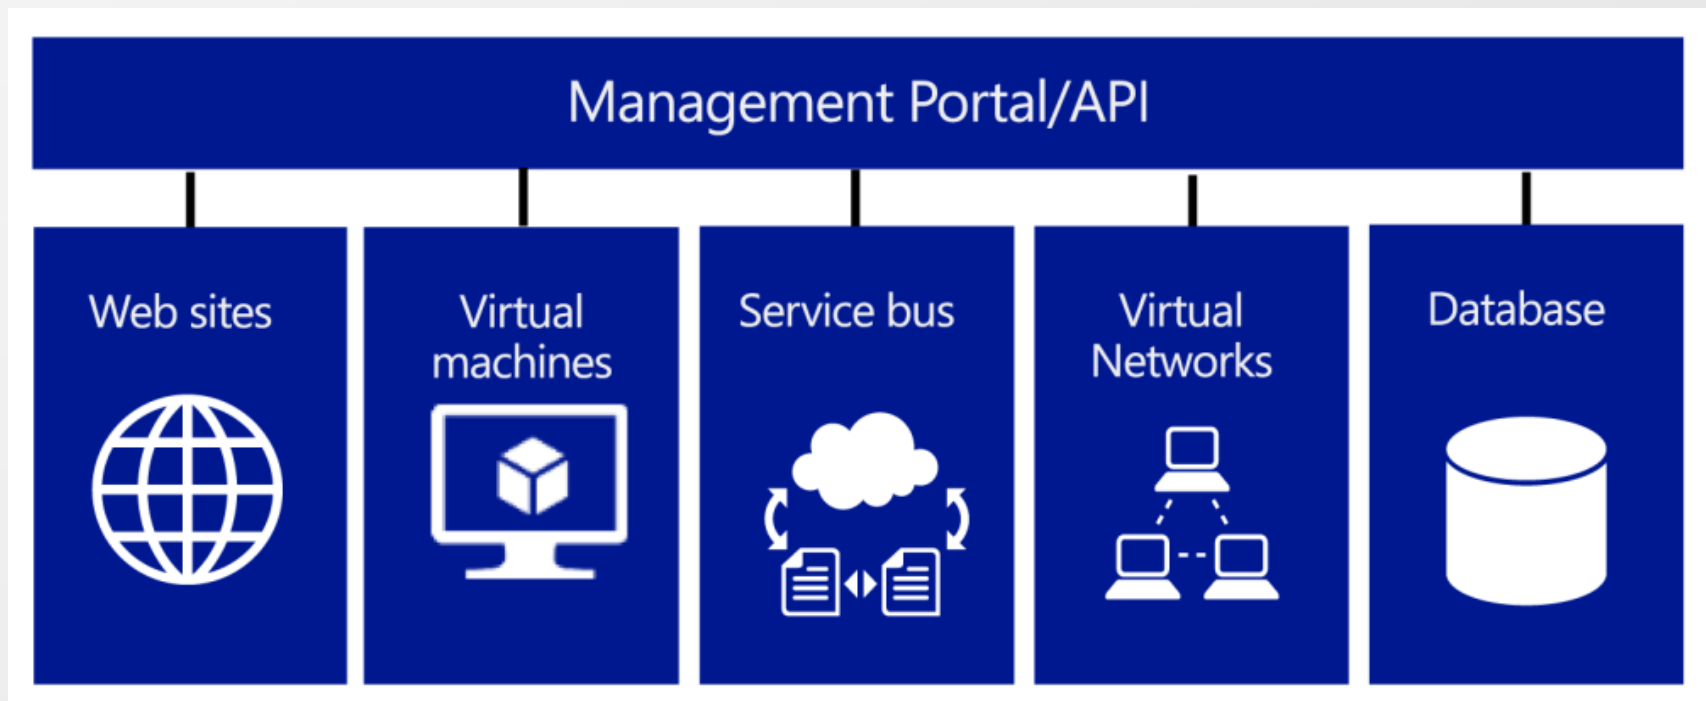
Windows Azure - Staging

- ❏ Allows to do application upgrade without downtime
- ❏ First new version is prepared into staging area
- ❏ You test the application in staging area
- ❏ Then staging area becomes production and vice versa
- ❏ Switch is done on the load balancer - it redirects all requests to the new version
- ❏ Old version is still running (currently in staging area). You are charged!



Windows Azure Pack

- ❑ Windows Azure technologies available in your data center
- ❑ Limited set of services from “full Azure”



Microsoft Azure - Pricing

- ❏ Prices are dependant on datacenter location (valid only for some services, some services are data center location independent)
- ❏ Different virtual server sizes have different prices
- ❏ Prices are based on usage (used storage, amount of send messages, amount of data transfers, etc.)
- ❏ Prices are dependant on the used service (web, VM, database, ...)
- ❏ You pay when your application is deployed even if it is not running!
- ❏ Pricing details:
 - ❏ <http://www.windowsazure.com/en-us/pricing/overview>
- ❏ Pricing calculator:
 - ❏ <http://www.windowsazure.com/en-us/pricing/calculator>

Microsoft Azure - Pricing Trial

New customers will get for one month (March 2015)

- ❏ 150 USD of Windows Azure Credits
- ❏ You can combine any services up to your credit

Google Cloud Platform Overview

- ❏ Google Cloud Platform is IaaS and PaaS
- ❏ Started in April 2008 with App Engine (web applications)
- ❏ Data center regions (March 2015)
 - ❏ USA: 6 regions
 - ❏ South America: 1 region
 - ❏ Europe: 4 region
 - ❏ Asia: 2 region
- ❏ Each region can have one or more zones
 - ❏ Zones represent isolated locations within a region
 - ❏ Zones have high-bandwidth, low-latency network connections to other zones in the same region



Google Cloud Platform Overview of Services I

☒ Compute

- ☒ Compute Engine
- ☒ App Engine
- ☒ Container Engine (in Alpha)

☒ Networking

- ☒ Load Balancing
- ☒ Interconnect
- ☒ DNS (in Beta)

☒ Storage

- ☒ Cloud SQL
- ☒ Cloud Storage
- ☒ Cloud Data Store

☒ Big Data

- ☒ BigQuery
- ☒ DataFlow (in Alpha)

Google Cloud Platform Overview of Services II

Services

- Translate API
- Prediction API
- Cloud Endpoints
- Pub/Sub (in Beta)

Management

- Deployment Manager
- Monitoring

SLA (monthly uptime):

- 99.95% for Compute Engine, App Engine, Cloud SQL, Load Balancing
- 99.9% for Cloud Storage, Prediction API, BigQuery
- No SLA found for Container Engine, Cloud Data Store*, Translate API, Data Flow, Interconnect, Cloud DNS, Pub/Sub



Google Cloud Platform - Compute and Networking

❏ Compute Engine

- ❏ Google's Infrastructure as a Service (IaaS)

❏ App Engine

- ❏ Run web applications in the Google infrastructure

❏ Container Engine

- ❏ Docker containers on Google Cloud Platform's virtual machines

❏ Load Balancing

- ❏ Load balancing in Google cloud

❏ Interconnect


- ❏ Private (non-internet) connection to the Google cloud

❏ DNS

- ❏ Production quality, high volume DNS serving

Google Cloud Platform - Storage and BigData

- ❏ **Cloud SQL**
 - ❏ MySQL database in Google cloud
- ❏ **Cloud Storage**
 - ❏ Key-value storage for big objects
- ❏ **Cloud Datastore**
 - ❏ NoSQL, schemaless database
- ❏ **Big Query**
 - ❏ Data analysis service
- ❏ **Data Flow**
 - ❏ Scalable data processing



Google Cloud Platform - Services and Management

❏ Cloud Endpoints

- ❏ RESTful services for mobile clients built on top of services in App Engine

❏ Translate API

- ❏ Machine translation between different languages

❏ Prediction API

- ❏ Pattern-matching and machine learning capabilities

❏ Pub/Sub

- ❏ Asynchronous messaging between applications

❏ Deployment Manager

- ❏ Create, deploy, and manage Google Cloud Platform resources

❏ Monitoring

- ❏ Monitoring service for your cloud applications in Google Cloud

Google Cloud Platform - App Engine I

Supported environments

❏ Java

- ❏ Java 7 JVM in safe sandboxed environment
- ❏ You can use JSP or JSF (JavaServer Pages/Faces)
- ❏ Some limits exist. E.g. you cannot
 - ❏ Write to the file system
 - ❏ Create threads in some ways
 - ❏ Directly communicate with system or native code via JNI
 - ❏ Create server sockets
- ❏ There are limits for web application in general e.g.
 - ❏ Maximal response time 60 seconds
 - ❏ Maximal request and response size 32MB



Google Cloud Platform - App Engine II

Supported environments (continue)

- ❏ Java
 - ❏ You can use languages compatible with JVM: JRuby, Scala, ...
- ❏ Python
 - ❏ Python 2.5 or 2.7
 - ❏ Similar limits like for Java environment
- ❏ Go (Beta in March 2015)
 - ❏ Go (1.2) is a new Google's language
 - ❏ Similar limits like for Java or Python environment
- ❏ PHP (Beta in March 2015)
 - ❏ PHP 5.4 in sandboxed environment
- ❏ Managed VMs (Beta in March 2015)
 - ❏ Virtual machines without sandbox constrains for Java, Python and Go

Google Cloud Platform - App Engine III

- ❏ Application is decomposed into modules
- ❏ Each module can have several versions
- ❏ Module version runs in some VM instance
- ❏ Several instance classes (with different memory and CPU)
- ❏ Different scaling possibilities
 - ❏ Manual Scaling
 - ❏ Runs continuously
 - ❏ Basic Scaling
 - ❏ Create an instance when the application receives a request
 - ❏ Automatic Scaling
 - ❏ Scaling based on different metrics

Google Cloud Platform - Storage - Datastore

- ❑ Schemaless NoSQL object data store with some features of “standard DBs” (e.g. transactions and references to other entities)
- ❑ Objects in data store are entities with one or more properties
- ❑ Properties have types like integers, strings, reference to entity, etc.
- ❑ Entities are of given kind (like person)
- ❑ Entities can be hierarchically structured
- ❑ For querying objects is used GQL (SQL-like syntax without “join”)
- ❑ Consistency
 - ❑ Strong consistency for reads and ancestor queries
 - ❑ Eventual consistency for other queries
- ❑ Maximum entity size is 1 MB

Google Cloud Platform - Storage - Cloud Storage I

- ❏ Key-value storage for big objects (“bucket” based)
- ❏ Consistency
 - ❏ Strong read after write consistency for objects (it is guaranteed that immediate read after write will return object)
 - ❏ Strong read after delete consistency for objects
 - ❏ “List” operations are eventually consistent
- ❏ Different versions of the same object possible
- ❏ Automatic lifecycle of objects (e.g. keep only last 3 versions)
- ❏ You can specify location of bucket: US, Europe or Asia
- ❏ Maximum file size is 5TB

Google Cloud Platform - Storage - Cloud Storage II

- ❑ Different storage classes are possible
 - ❑ Standard storage
 - ❑ Durable reduced availability
 - ❑ Nearline storage
- ❑ Data accessible as
 - ❑ <https://storage.cloud.google.com/bucket/object>
- ❑ Redirects via DNS CNAME possible - <http://yourdomain/object>
- ❑ Supports TLS (HTTPS)
- ❑ Support access control (different for buckets and objects)
- ❑ Supports server-side encryption (data are written encrypted to the disk)
- ❑ Supports resumable uploads
- ❑ Supports streaming

Google Cloud Platform - Storage - Cloud SQL

- ❑ Based on MySQL 5.5 or 5.6
- ❑ Data stored in US, EU or Asia region
- ❑ Maximal database size is 500GB (250GB by default, higher size on demand)
- ❑ All database requests must finish within 60 seconds (frontend request) or 10 minutes (offline request)
- ❑ Each App Engine cannot have more than 12 concurrent connections to SQL
- ❑ Data are internally encrypted
- ❑ Some features are not supported against standard MySQL
 - ❑ User defined functions
 - ❑ MySql Replication
 - ❑ Some MySQL statements and functions

Google Cloud Platform - Pricing

Google Cloud Pricing Calculator

- <https://cloud.google.com/products/calculator>

Pricing details

- <https://developers.google.com/compute/pricing>

- <https://developers.google.com/appengine/pricing>

- <https://developers.google.com/cloud-sql/pricing>

- <https://developers.google.com/storage/pricing>

- <https://developers.google.com/bigquery/pricing>

- <https://developers.google.com/prediction/pricing>

- ...

- Free trial quota is available as well

Pros of cloud computing

- ❏ You don't need to maintain your infrastructure
 - ❏ You don't need to have unused (reserved) hardware
 - ❏ You don't need appropriate specialists
- ❏ Pay only for what you use
- ❏ Costs can be lower
 - ❏ Especially in cases when you expect usage peaks
- ❏ Easy to scale for increased number of users
 - ❏ But application must be prepared as well

Cons of cloud computing

- ❏ Security and privacy
- ❏ Strong dependency on the network availability
- ❏ Vendor lock-in
 - ❏ There aren't standards to migrate between vendors
- ❏ Migration costs from existing infrastructure
- ❏ Legal aspects
 - ❏ Provider and user can have different law
- ❏ Cloud platforms provide less functionality than existing platforms
- ❏ Requires new know-how for the whole team

Děkuji za pozornost.

Tento projekt je spolufinancován Evropským sociálním fondem a státním rozpočtem České republiky.



INVESTICE DO ROZVOJE VZDĚLÁVÁNÍ