



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.



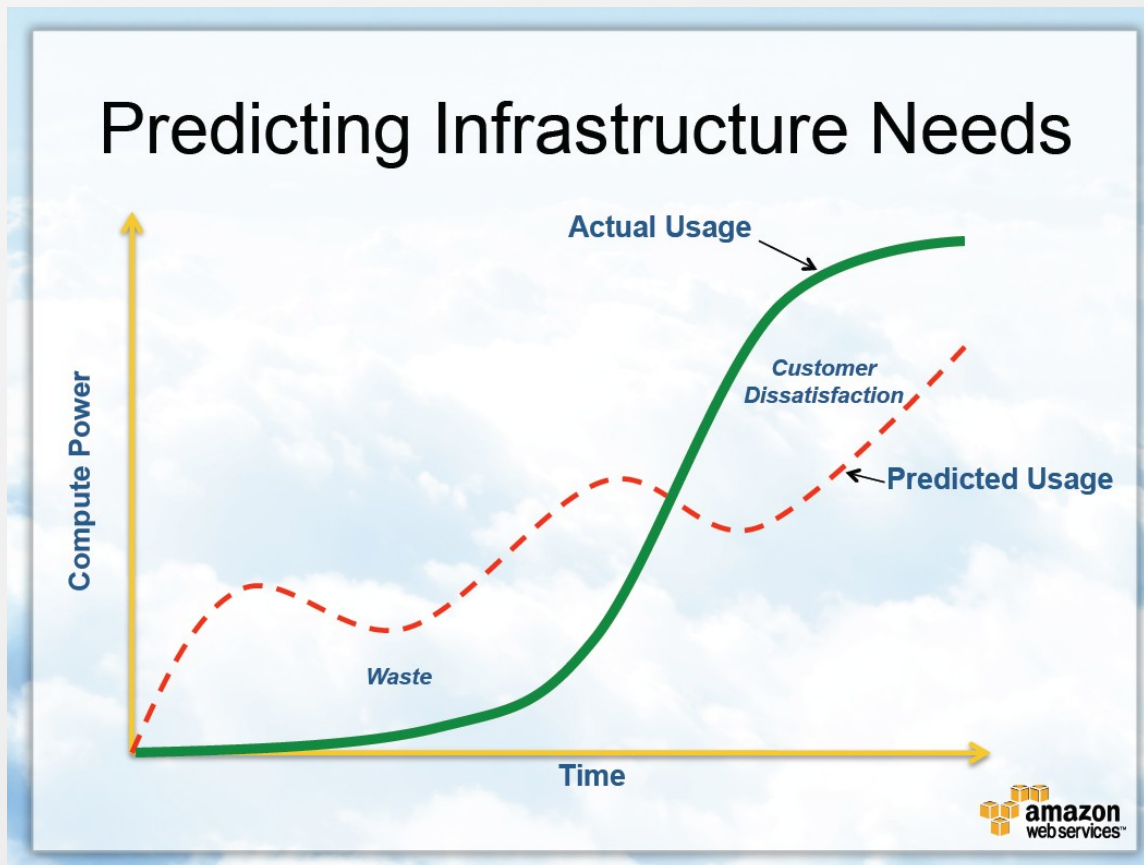
INVESTICE DO ROZVOJE VZDĚLÁVÁNÍ

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.

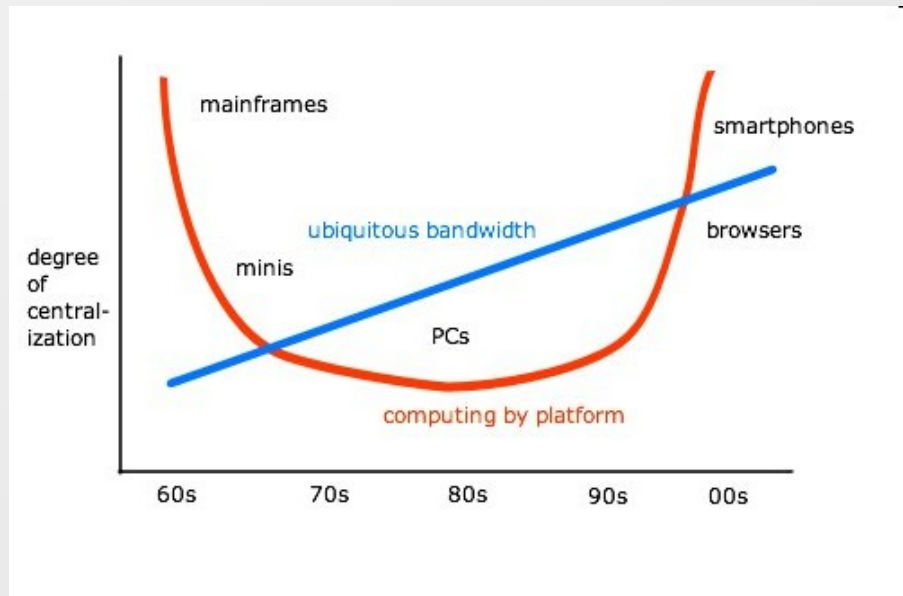
Five Attributes of Cloud Computing

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




History

- 1960s First ideas
- 1999 Salesforces.com
- 2002 Google Search API
Amazon Web Services
- 2006 Amazon Elastic Cloud
- 2008 Google App Engine
Microsoft Azure Platform
- 2009 VMware vCloud Express
- 2010 Microsoft Office WebApps
Google Services



Cloud Services Stack I



Software
as a Service

Platform
as a Service

Infrastructure
as a Service

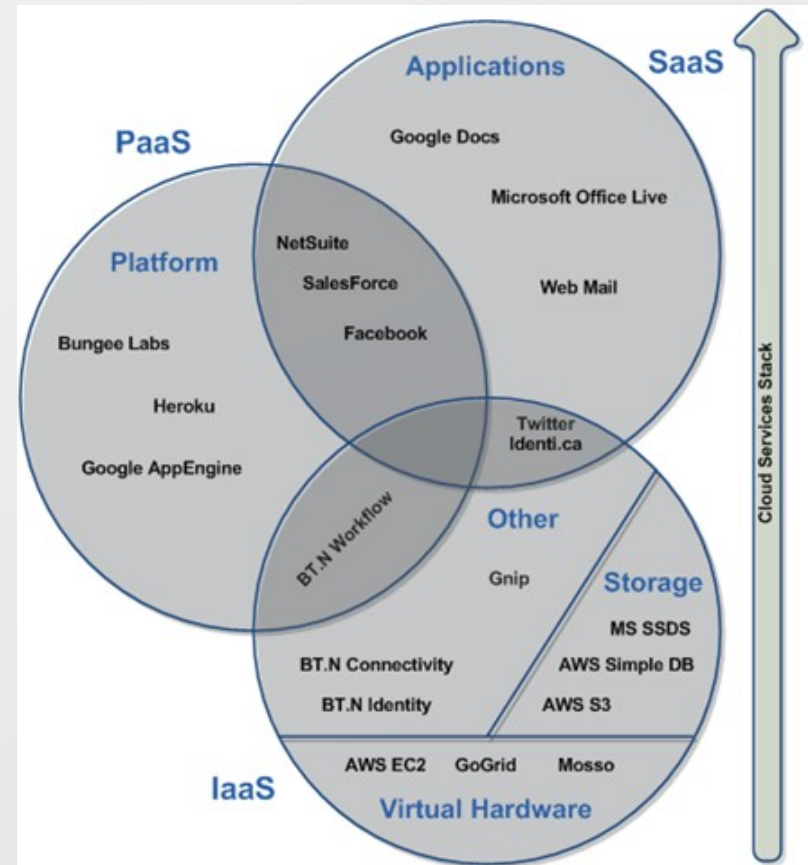
Cloud Services Stack II

- **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


Cloud Services Stack III

Some types of services overlap

- ❑ Borders are “foggy”



Cloud Services Stack - Examples

Cloud Market Types	Types of Offerings	Examples	Level of Abstraction 
Software-as-a-Service	<ul style="list-style-type: none"> • Rich Internet application web sites • Application as Web Sites • Collaboration and email • Office Productivity • Client apps that connect to services in the cloud 	<ul style="list-style-type: none"> • Flickr • Myspace.com • Cisco WebEX office • Gmail • IBM Bluehouse 	
App-components-as-a-Service	<ul style="list-style-type: none"> • APIs for specific service access for integration • Web-based software service than can combine to create new services, as in a mashup 	<ul style="list-style-type: none"> • Amazon Flexible Payments Service and DevPay • Salesforce.com's AppExchange • Yahoo! Maps API • Google Calendar API • zembly 	
Platform-as-a-Service	<ul style="list-style-type: none"> • Development-platform-as-a-service • Database • Message Queue • App Servicer • Blob or object data stores 	<ul style="list-style-type: none"> • Google App Engine and BigTable • Microsoft SQL Server Data Services • Engine Yard • Salesforce.com's Force.com 	
Infrastructure-as-a-Service	<ul style="list-style-type: none"> • Virtual servers • Logical disks • VLAN networks • Systems Management 	<ul style="list-style-type: none"> • Akamai • Amazon EC2 • CohesiveFT • Mosso (from Rackspace) • Joyent Accelerators • Nirvanix Storagee Delivery Network 	
Physical Infrastructure	<ul style="list-style-type: none"> • Managed Hosting • Collocation • Internet Service Provider • Unmanaged hosting 	<ul style="list-style-type: none"> • GoDaddy.com • Rackspace • Savvis 	


Adapted from Forrester Research Taxonomy

Types of cloud platforms

- ❏ **Public cloud**
 - ❏ “Classical” model
 - ❏ The same functionality for all users
- ❏ **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



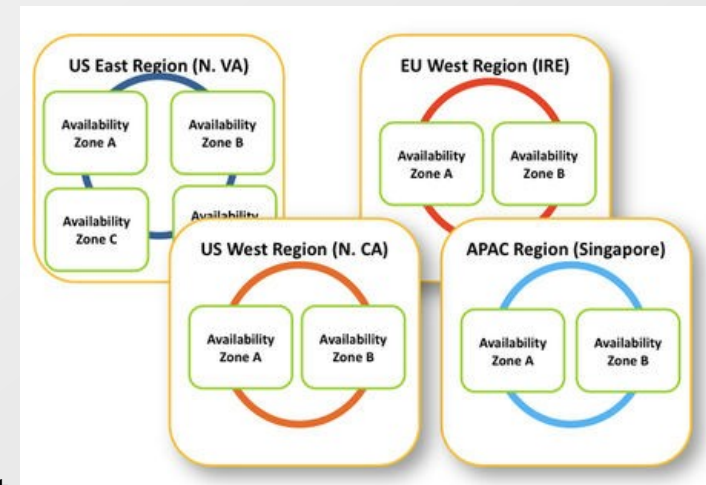
Amazon is a leader in IaaS and PaaS
AWS introduced in 2002

- 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 using REST and SOAP
- Number of services increase every year

Amazon Web Services - Regions and availability zones I

Available regions (February 2012)

- ❑ US East - Northern Virginia (December 2008)
- ❑ US West - Oregon (November 2011)
- ❑ US West - Northern California (December 2009)
- ❑ EU - Ireland (December 2008)
- ❑ Asia-Pacific - Singapore (April 2010)
- ❑ Asia-Pacific - Tokyo (April 2011)
- ❑ South America - Sao Paulo (December 2011)
- ❑ AWS GovCloud - Special region for US government




Availability zones

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

Amazon Web Services - Regions and availability zones II

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



Amazon Web Services - Portfolio I

Computing

- ❑ Elastic Compute Cloud (EC2, August 2006)
 - ❑ Virtualized OS in Amazon's infrastructure
- ❑ Elastic MapReduce (April 2009)
 - ❑ Process vast amount of data
- ❑ Auto Scaling (May 2009)
 - ❑ Allows to automatically scale EC2 capacity up or down
- ❑ Elastic Load Balancing (May 2009)
 - ❑ Distributes requests to multiple EC2 instances

Content Delivery

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

Database

- ❑ Relation Database Service (October 2009)
 - ❑ Relational database in the cloud (MySQL, Oracle)
- ❑ DynamoDB (January 2012)
 - ❑ NoSQL database



Amazon Web Services - Portfolio II

Database

- ❏ SimpleDB (December 2007)
 - ❏ Non-relational data store
- ❏ ElastiCache (August 2011)
 - ❏ In-memory cache in the cloud

Deployment and Management

- ❏ Identity and Access Management (IAM, September 2010)
 - ❏ Securely controls access to services and resources
- ❏ CloudWatch (May 2009)
 - ❏ Monitoring of cloud resources
- ❏ Elastic Beanstalk (January 2011)
 - ❏ Deployment and management tool
- ❏ CloudFormation (February 2011)
 - ❏ Tool for creating templates of related AWS resources

Amazon Web Services - Portfolio III

Messaging

- ❏ Simple Queue Service (July 2006)
 - ❏ Message queue in the cloud
- ❏ Simple Notification Service (April 2010)
 - ❏ Sends notifications from the cloud
- ❏ Simple Email Service (January 2011)
 - ❏ Scalable e-mail sending service

Networking

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

Amazon Web Services - Portfolio IV

Payments & Billing

- ❏ Flexible Payments Service (August 2007)
 - ❏ Payment service for developers
- ❏ DevPay (December 2007)
 - ❏ Simple to use payment and account management service

Storage

- ❏ Simple Storage Service (S3, March 2006)
 - ❏ Key-value storage for big objects
- ❏ Elastic Block Store (EBS, August 2008)
 - ❏ Block device mounted by EC2 instances
- ❏ Import/Export (May 2009)
 - ❏ Allows to import/export huge amount of data from/to portable device
- ❏ Storage Gateway (January 2012)
 - ❏ Backup service in the cloud

Amazon Web Services - Portfolio V

Web Traffic

- ❏ Alexa Web Information Service (October 2004)
 - ❏ Web statistics
- ❏ Alexa Top Sites (January 2006)
 - ❏ Top internet sitest statistics

Workforce

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



Amazon Web Services - Compute Services I

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

- ❑ Standard small (default) - 1.7 GB memory, 1 virtual core, 1 ECU, 160 GB instance storage, 32 bit
- ❑ Micro - 613 MB memory, up to 2 ECUs, EBS storage only, 32 or 64 bit
- ❑ High-Memory Quadruple Extra Large - 68.4 GB memory, 26 ECUs (8 virtual cores with 3.25 ECUs each), 1690 GB instance storage, 64 bit
- ❑ High-CPU Extra Large - 7GB memory, 20 ECUs (8 virtual cores with 2.5 ECUs each), 1690 GB instance storage, 64 bit
- ❑ Cluster compute Eight Extra Large - 60.5 GB memory, 88 ECUs, 3370 instance storage, 64 bit, 10 Gbit ethernet
- ❑ Cluster GPU Quadruple Extra Large - 22 GB memory, 33.5 ECUs, 2x NVIDIA Tesla “Fermi” M2050 GPUs, 1690 instance storage, 64 bit, 10 Gbit ethernet
- ❑ Small DB Instance: 1.7 GB memory, 1 virtual core, 1 ECU, 64-bit platform, Moderate I/O Capacity
- ❑ High-Memory Quadruple Extra Large DB Instance: 68 GB of memory, 26 ECUs (8 virtual cores with 3.25 ECUs each), 64-bit platform, High I/O Capacity

ECU = EC2 Compute Unit. Equivalent of 1.0-1.2 GHz 2007 Opteron or Xeon processor

Amazon Web Services - Compute Services II

Supported operating systems

- Linux: Red Hat Enterprise, SUSE Enterprise, Oracle Enterprise, Amazon AMI, Ubuntu, Fedora, Gentoo, Debian
- Windows server: 2003 R2, 2008, 2008 R2
- You can also use your own OS

Pre-configured software

- Databases: Oracle 11g, MySQL Enterprise, Microsoft SQL Server, IBM DB2, IBM Informix Dynamic Server
- Web hosting: Apache, IIS/ASP.NET, IBM WebSphere Portal Server, IBM Lotus Web Content Management
- Resource management: Hadoop, Condor, StackIQ Rocks+
- Application development environments: IBM sMash, JBoss Enterprise Application Server, Ruby on Rails
- Application servers: IBM WebSphere Application Server, Java Application Server, Oracle WebLogic Server
- Video encoding & streaming: Wowza Media Server Pro, Windows Media Server

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

Amazon Web Services - Compute Services III

- ☒ Amazon Elastic 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
- ☒ 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 (file system) - can be attached just to one EC2 instance
 - ☒ For storage you have to pay extra



Amazon Web Services - Simple Storage Service (S3)

- ❏ 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)
- ❏ Supports HTTP, BitTorrent
- ❏ 99.999999999% durability, 99.99% availability
- ❏ With Reduced Redundancy Storage durability and availability 99.99%
- ❏ “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

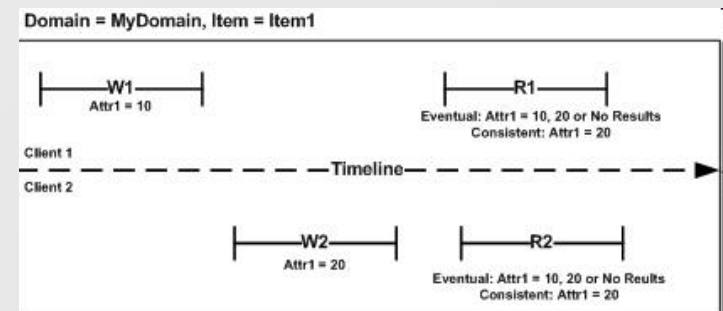
Amazon Web Services - Database Services I

Relational Database Service

- ☒ Allows to easily migrate your existing applications
- ☒ Support for MySQL or Oracle
- ☒ By default you are allowed to run up to 20 instances (but you can ask for more)
- ☒ Size limit per instance is 1TB

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)
- ☒ 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 II


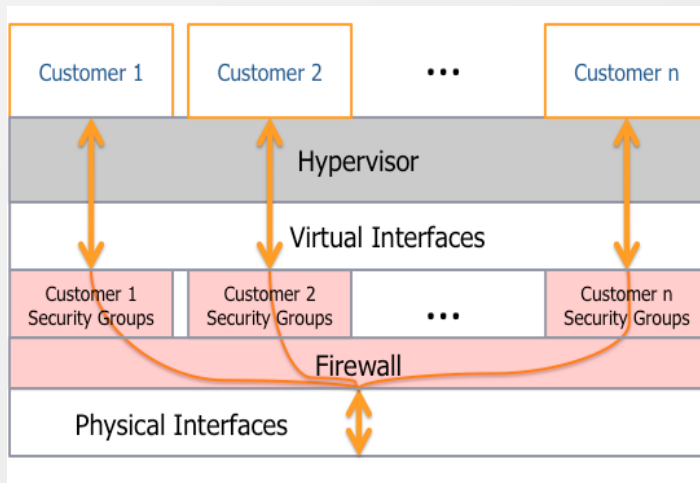
☒ DynamoDB

- ☒ NoSQL (non-relational) database
- ☒ Fast predictable performance, stored all data on SSD
- ☒ No limits on amount of data, scales automatically
- ☒ Table - item - attribute approach
- ☒ Requires primary key for item, only primary key is indexed
- ☒ Primary key can be compound (hash and range type primary key)
- ☒ Item size limit is 64kB (sum of all attribute names and values)
- ☒ Not so flexible queries
- ☒ Consistency
 - ☒ Eventually consistent reads
 - ☒ Consistent reads
- ☒ Integrates Elastic MapReduce

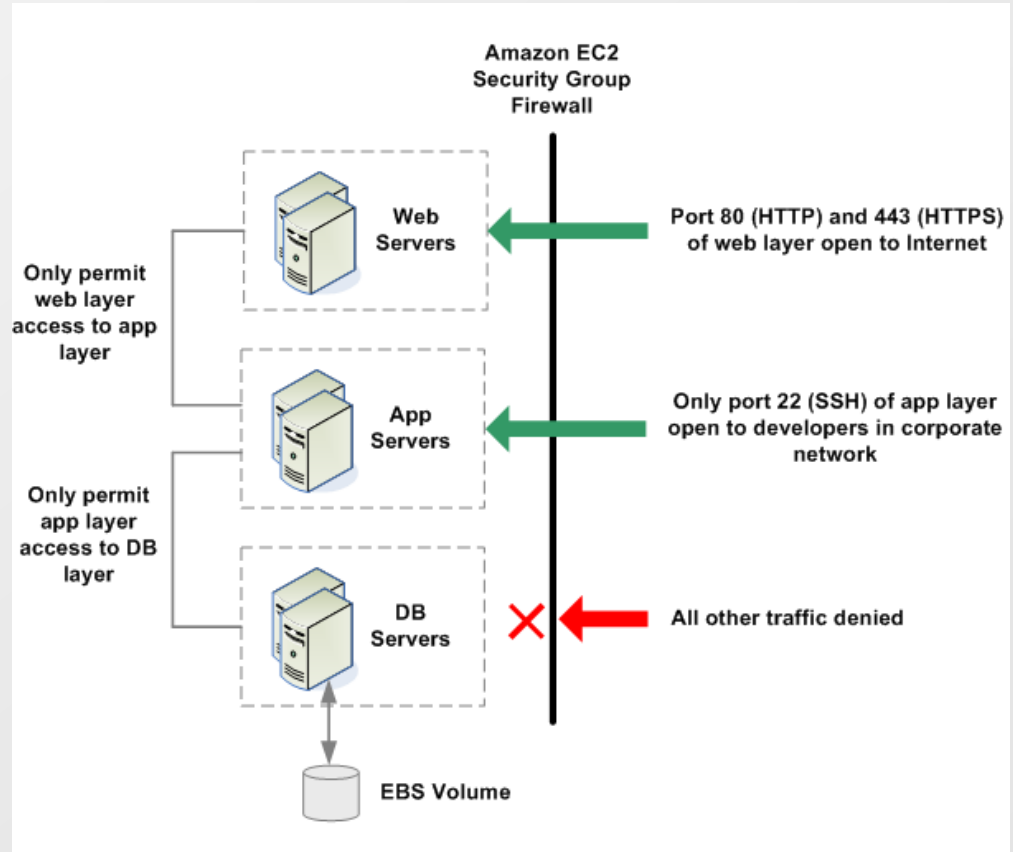
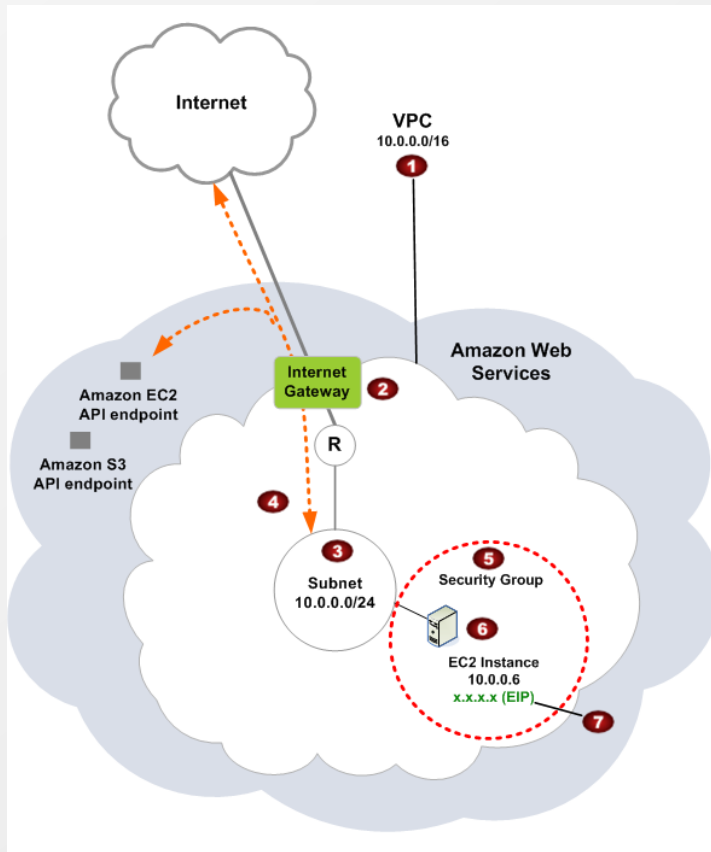
You can run also other databases like Microsoft SQL Server, IBM DB2, Postgre SQL, Sybase, etc. but they are not officially supported by AWS.

Amazon Web Services - Virtual Private Cloud I

- ☒ Connection via VPN
- ☒ Isolated network access
- ☒ Subnets

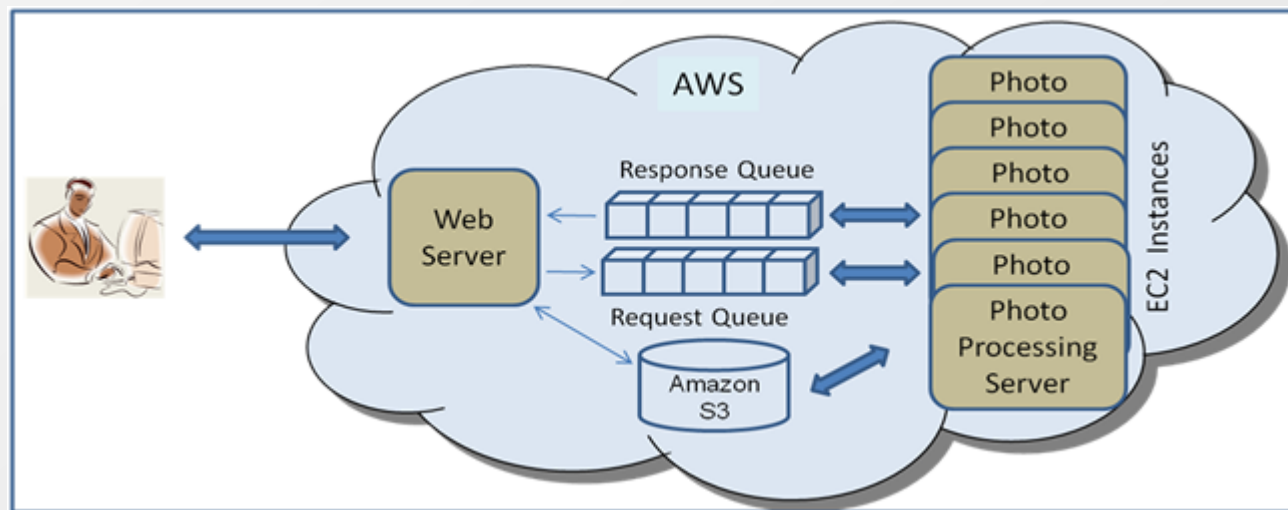


Amazon Web Services - Virtual Private Cloud, Security



Amazon Web Services - EC2 and S3 Example

Photo processing example



Amazon Web Services - Pricing

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 EC2 I

Pricing for EC2 on-demand instances (EU Region, February 2012)

EC2 instance	Linux per hour	Windows per hour
Standard Small	\$0.095	\$0.12
Standard Large	\$0.38	\$0.48
Standard Extra Large	\$0.76	\$0.96
Micro	\$0.025	\$0.035
Hi-Memory Extra Large	\$0.57	\$0.62
Hi-Memory Double Extra Large	\$1.14	\$1.24
Hi-Memory Quadruple Extra Large	\$2.28	\$2.48
Hi-CPU Medium	\$0.19	\$0.29
Hi-CPU Extra Large	\$0.76	\$1.16
Cluster Compute Quadruple Extra Large	N/A	N/A
Cluster GPU Quadruple Extra Large	N/A	N/A

Amazon Web Services - Pricing EC2 II

Pricing for EC2 light utilization reserved instances (EU Region, Feb. 2012)

EC2 instance	1 year term	3 year term	Linux per hour	Windows per hour
Standard Small	\$97.50	\$150	\$0.063	\$0.082
Standard Large	\$390	\$600	\$0.25	\$0.33
Standard Extra Large	\$780	\$1200	\$0.50	\$0.66
Micro	\$23	\$35	\$0.015	\$0.021
Hi-Memory Extra Large	\$555	\$852.50	\$0.375	\$0.455
Hi-Memory Double Extra Large	\$1100	\$1705	\$0.75	\$0.91
Hi-Memory Quadruple Extra Large	\$2200	\$3410	\$1.50	\$1.82
Hi-CPU Medium	\$195	\$300	\$0.125	\$0.19
Hi-CPU Extra Large	\$780	\$1200	\$0.50	\$0.76
Cluster Compute Quadruple Extra Large	N/A	N/A	N/A	N/A
Cluster GPU Quadruple Extra Large	N/A	N/A	N/A	N/A

Amazon Web Services - Pricing EC2 III

Pricing for EC2 medium utilization reserved instances (EU Region, Feb. 2012)

EC2 instance	1 year term	3 year term	Linux per hour	Windows per hour
Standard Small	\$227.50	\$350	\$0.04	\$0.06
Standard Large	\$910	\$1400	\$0.16	\$0.24
Standard Extra Large	\$1820	\$2800	\$0.32	\$0.48
Micro	\$54	\$82	\$0.01	\$0.016
Hi-Memory Extra Large	\$1325	\$2000	\$0.24	\$0.32
Hi-Memory Double Extra Large	\$2650	\$4000	\$0.48	\$0.64
Hi-Memory Quadruple Extra Large	\$5300	\$8000	\$0.96	\$1.28
Hi-CPU Medium	\$455	\$700	\$0.08	\$0.145
Hi-CPU Extra Large	\$1820	\$2800	\$0.32	\$0.68
Cluster Compute Quadruple Extra Large	N/A	N/A	N/A	N/A
Cluster GPU Quadruple Extra Large	N/A	N/A	N/A	N/A

Amazon Web Services - Pricing EC2 IV

Pricing for EC2 heavy utilization reserved instances (EU Region, Feb. 2012)

EC2 instance	1 year term	3 year term	Linux per hour	Windows per hour
Standard Small	\$227.50	\$425	\$0.031	\$0.051
Standard Large	\$1105	\$1700	\$0.125	\$0.205
Standard Extra Large	\$2210	\$3400	\$0.25	\$0.41
Micro	\$62	\$100	\$0.008	\$0.014
Hi-Memory Extra Large	\$1600	\$2415	\$0.19	\$0.27
Hi-Memory Double Extra Large	\$3200	\$4830	\$0.38	\$0.54
Hi-Memory Quadruple Extra Large	\$6400	\$9660	\$0.76	\$1.08
Hi-CPU Medium	\$553	\$850	\$0.063	\$0.128
Hi-CPU Extra Large	\$2210	\$3400	\$0.25	\$0.51
Cluster Compute Quadruple Extra Large	N/A	N/A	N/A	N/A
Cluster GPU Quadruple Extra Large	N/A	N/A	N/A	N/A

Amazon Web Services - Pricing EC2 V

Lowest pricing for EC2 spot instances (EU Region, Feb. 2012)

EC2 instance	Linux per hour	Windows per hour
Standard Small	\$0.036	\$0.06
Standard Large	\$0.144	\$0.24
Standard Extra Large	\$0.288	\$0.48
Micro	\$0.009	\$0.014
Hi-Memory Extra Large	\$0.216	\$0.288
Hi-Memory Double Extra Large	\$0.504	\$0.66
Hi-Memory Quadruple Extra Large	\$1.008	\$1.32
Hi-CPU Medium	\$0.072	\$0.15
Hi-CPU Extra Large	\$0.288	\$0.6
Cluster Compute Quadruple Extra Large	N/A	N/A
Cluster GPU Quadruple Extra Large	N/A	N/A

Amazon Web Services - Pricing EC2 VI

Pricing for EC2 data transfers (EU Region, February 2012)

Type of transfer	Price per GB
Data Transfer In	Free
Data transfer out - first 1 GB per month	Free
Data transfer out - up to 10 TB per month	\$0.12
Data transfer out - next 40 TB per month	\$0.09
Data transfer out - next 100 TB per month	\$0.07
Data transfer out - next 350 TB per month	\$0.05

Pricing for EC2 data transfers (EU Region, May 2010) - history

Type of transfer	Price per GB	Comment
Data Transfer In	Free	\$0.10 after Jun 2010
Data transfer out - first 1 GB per month	Free	
Data transfer out - up to 10 TB per month	\$0.15	
Data transfer out - over 150 TB per month	\$0.08	

Amazon Web Services - Pricing EC2 VII

New AWS customers get each month for one year (February 2012)

- ❏ 750 hours of EC2 running Linux/Unix Micro instance usage
- ❏ 750 hours of EC2 running Microsoft Windows Server 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
- ❏ 15 GB of bandwidth out aggregated across all AWS services
- ❏ 1 GB of Regional Data Transfer

Amazon Web Services - Pricing S3

Pricing for S3 (EU Region, February 2012)

Storage per GB per month	Data transfers per GB per month	Requests
First 1 TB - \$0.125	First 1GB out - free	PUT - \$0.01 per 1000 requests
Over 5000 TB - \$0.055	Over 150 TB out - \$0.050	GET - \$0.01 per 10000 requests

Amazon Pricing Pages

- <http://aws.amazon.com/s3/pricing/>
- <http://aws.amazon.com/ec2/pricing/>
- ...

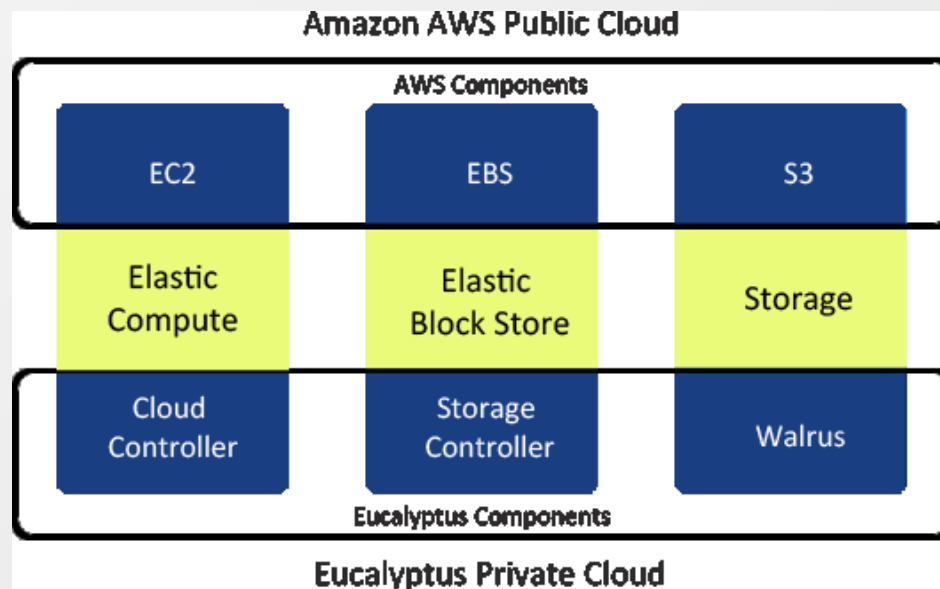
Simple Monthly Calculator

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

Amazon Web Services - Alternatives

Open source alternative - Eucalyptus

- ☞ Runs some of Amazon Web Services under your Linux
- ☞ <http://www.eucalyptus.com>



Microsoft Azure Platform

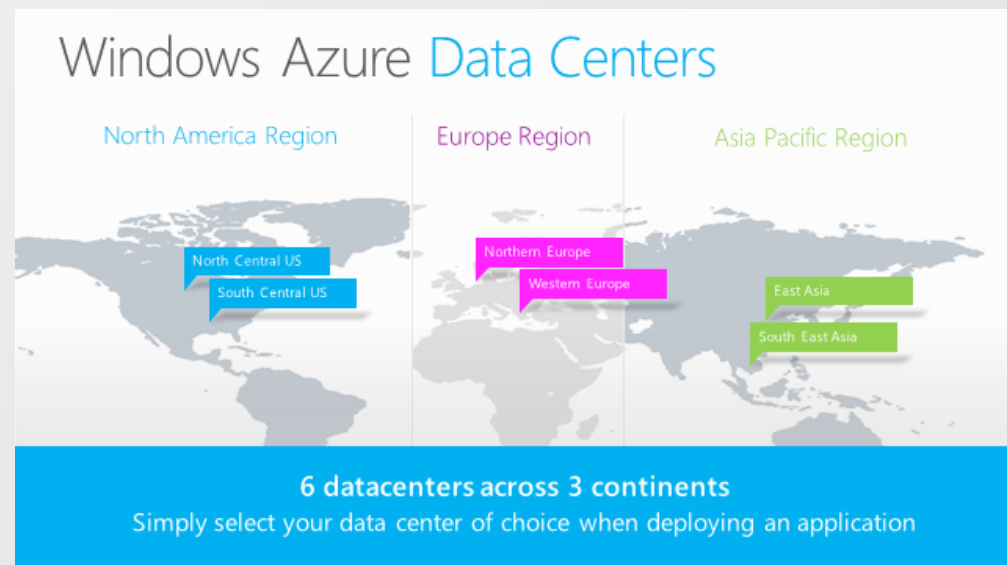
- ❏ Microsoft's response to the cloud computing
 - ❏ Platform and Software as a Service
- ❏ Runs on Windows Azure OS
 - ❏ Manages computing and storage resources
 - ❏ Windows Server (2008 R2 in February 2012) used for running applications
- ❏ Offers several services
- ❏ Available in several regions
- ❏ SLAs (monthly)
 - ❏ 99.95% for compute
 - ❏ 99.8% for storage, CDN, SQL Azure, Service Bus, Access Control



Windows Azure Data Centers

Data Centers (February 2012)

- ❑ US North Central - Chicago, IL
- ❑ US South Central - San Antonio, TX
- ❑ West Europe - Amsterdam
- ❑ North Europe - Dublin
- ❑ East Asia - Hong Kong
- ❑ South-East Asia - Singapore



Windows Azure - Services I

Windows Azure Compute

- ▣ Web Role
 - ▣ Runs web application
- ▣ Worker Role
 - ▣ Runs background processing
- ▣ VM Role
 - ▣ Customer's modified image of Windows Server 2008 R2 (February 2012)

Windows Azure Storage

- ▣ BLOB (Binary Large Object)
 - ▣ Key-value storage for big data
- ▣ Table
 - ▣ NoSQL database
- ▣ Queue
 - ▣ Message queue in the Microsoft's cloud
- ▣ Azure Drive
 - ▣ Standard NTFS drive which can be mounted to compute instance

Windows Azure - Services II

SQL Azure

- SQL Azure Database
 - Microsoft SQL Server in the cloud
- SQL Azure Data Sync
 - Synchronizes SQL databases (either Azure <-> On premises or Azure <-> Azure)
- SQL Azure Reporting
 - Reporting services on top of SQL Azure Database

Content Delivery Network (CDN)

- Delivery of static content or streaming media (24 locations worldwide, February 2012)

Service Bus

- Service Bus Messaging
 - Message queue in the cloud (more powerful than storage queue, newer)
- Service Bus Connectivity
 - Connects existing applications outside cloud

Windows Azure - Services III

Azure Active Directory

- ☒ Identity and Access management (IAM) in the cloud

Caching

- ☒ Distributed, in-memory cache

Azure Virtual Network

- ☒ Azure Connect
 - ☒ Allows to integrate existing applications with cloud applications
- ☒ Azure Traffic Manager
 - ☒ Load ballancer

HPC (High Performance Computing) Scheduler

- ☒ Allows to easily migrate compute-intensive applications into cloud

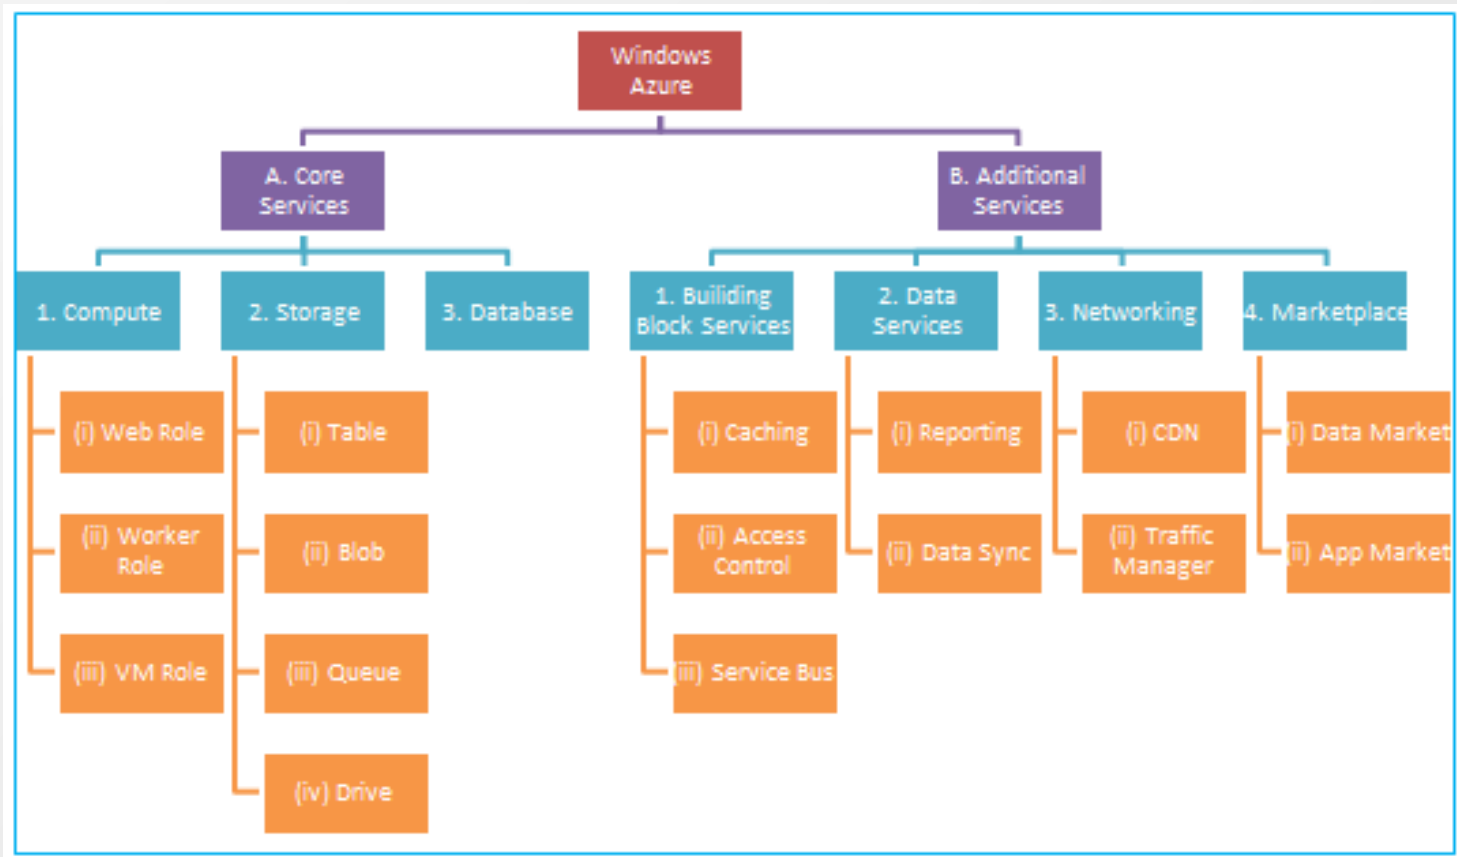
Azure Market Place

- ☒ Market place for selling applications and services

SendGrid

- ☒ Email service in the cloud

Windows Azure - Services III



Microsoft Azure - Computation

Web role

- ❏ Contains IIS
- ❏ Runs web applications

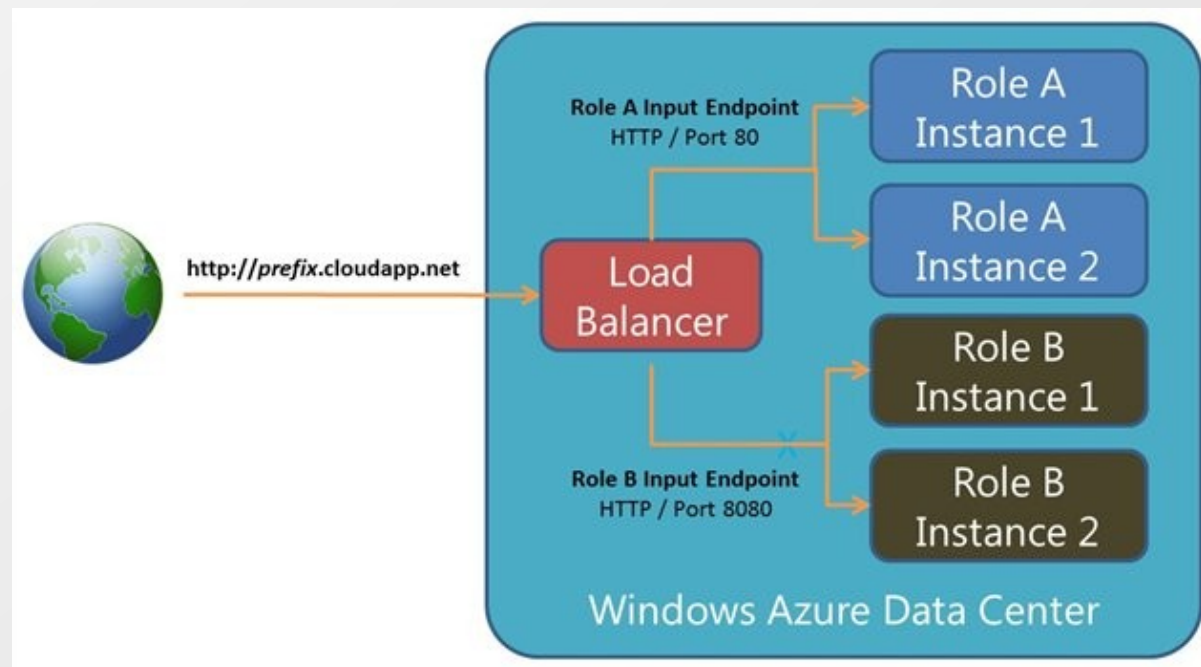
Worker role

- ❏ Runs tasks in the background

VM role

- ❏ Customized instance of Windows Server 2008 R2

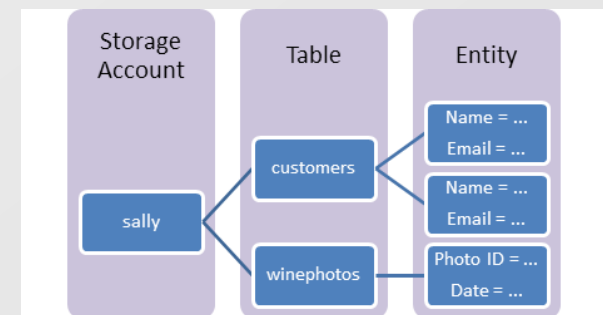
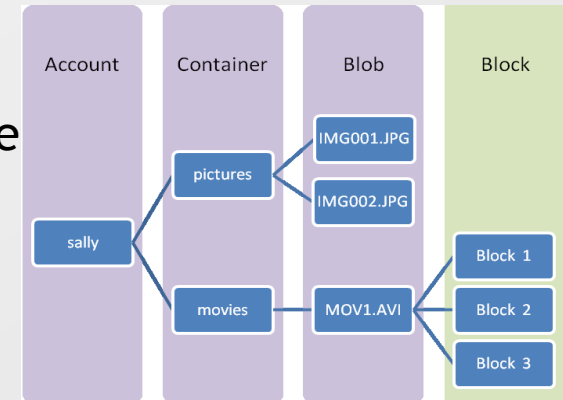
Web and worker role can communicate via WCF or queues



Microsoft Azure - Simple Data Storage

Simple Data Storage

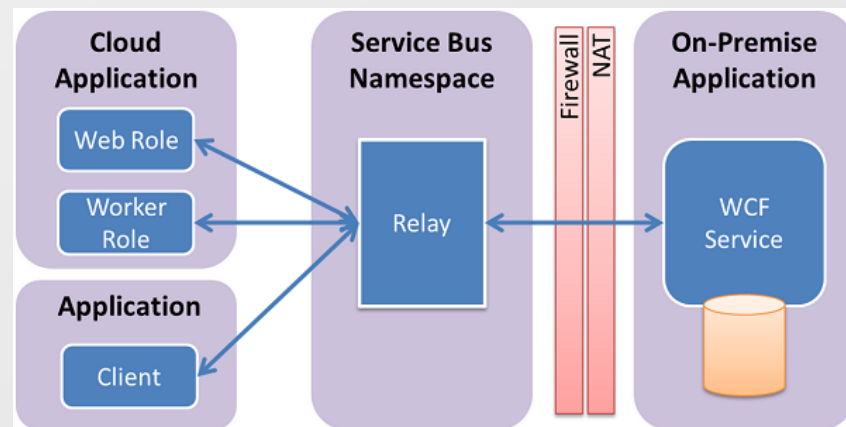
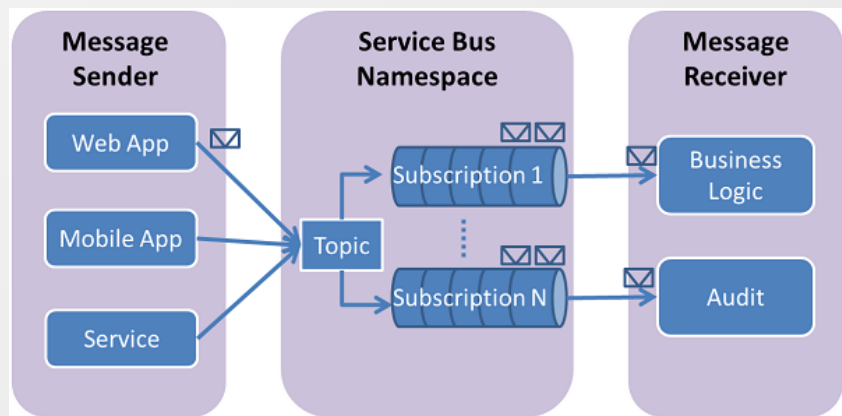

- ❏ Maximal 100TB of data per account for all storage
- ❏ Blob - Key-value storage with “block” support
 - ❏ Up to 200GB for block blob
 - ❏ Up to 1TB for page blob
 - ❏ Optimized for random read/write
 - ❏ `http://<storage account>.blob.core.windows.net/<container>/<blob>`
- ❏ Table
 - ❏ NoSQL indexed database
 - ❏ Maximal entity size is 1MB
 - ❏ Entity can have up to 252 custom properties
- ❏ Queue
 - ❏ Maximal 64kB per message



Microsoft Azure - SQL Azure Database

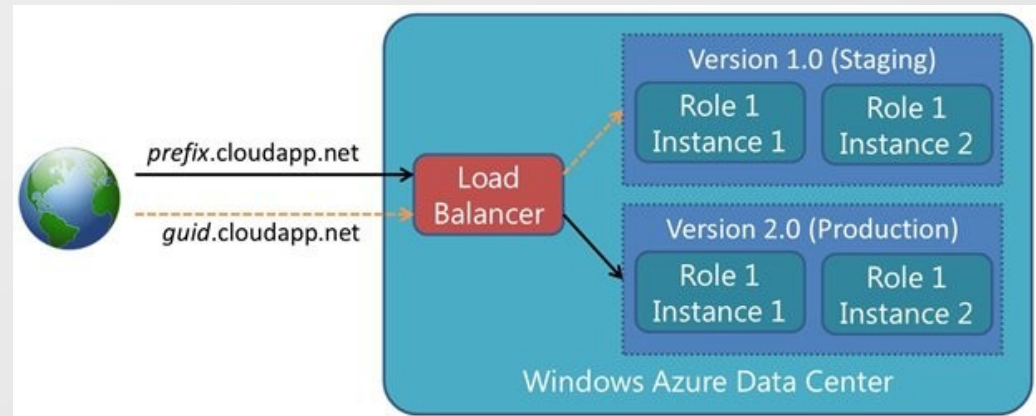

- Supports T-SQL
- Maximal 150 databases
- Maximal limits per database
 - Web Edition - 5GB
 - Business Edition - 150GB
- For larger data you have to use sharding (split data into several databases)
- No full text search, CLR, replication, database mirroring, jobs, distributed transactions, distributed queries

Microsoft Azure - Service Bus



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!



Microsoft Azure - Pricing Compute

Azure Compute Prices (February 2012):

Virtual Machine Size	CPU cores	Memory	Cost Per Hour
Extra Small	Shared	768 MB	\$0.04
Small	1	1.75 GB	\$0.12
Medium	2	3.5 GB	\$0.24
Large	4	7 GB	\$0.48
Extra Large	8	14 GB	\$0.96

You pay when your application is deployed even if it is not running!

Pricing calculator:

<http://www.windowsazure.com/en-us/pricing/calculator>

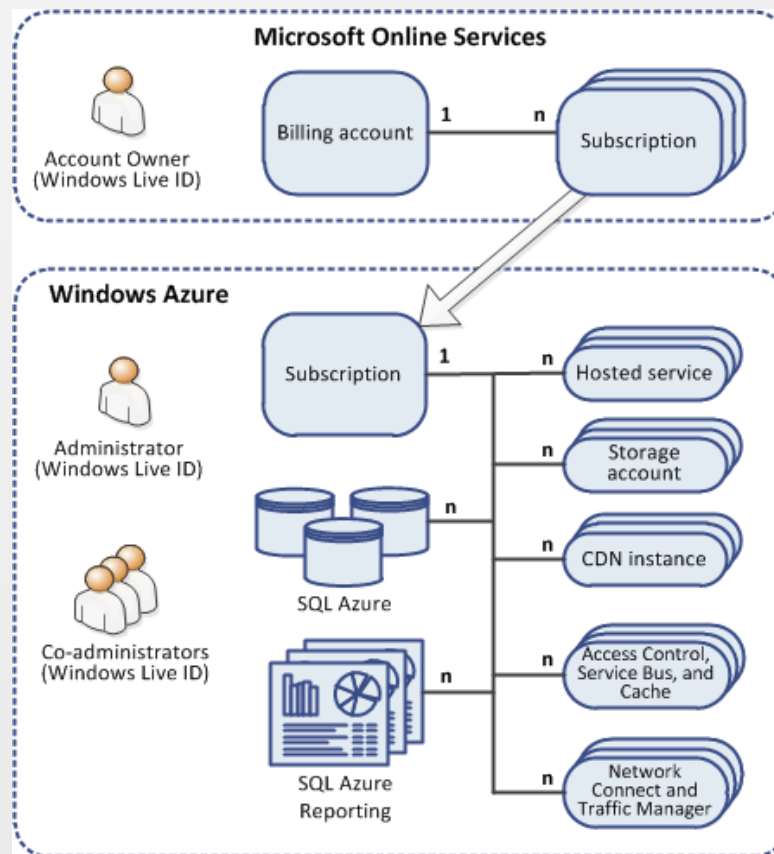
Microsoft Azure - Pricing Trial

New customers will get each month for three months (February 2012):

- ❏ Compute VM: 750 small compute hours per month
- ❏ Relational database: 1GB web edition SQL Azure database
- ❏ Storage: 20GB with 1,000,000 storage transactions
- ❏ Content delivery network: 500,000 CDN transactions
- ❏ Bandwidth: Unlimited inbound / 20GB Outbound

You need a credit card for signing into the trial.

Microsoft Azure - Billing Explanation



Google Cloud Services

- ☒ Google concentrates mainly on SaaS (Software as a Service)

- ☒ Gmail
- ☒ Calendar
- ☒ Documents
- ☒ CloudConnect



- ☒ Collaborative, multi-person editing for Microsoft Office applications
- ☒ Video
 - ☒ Communication channel for internal/external use

Google Cloud Services - PaaS

For own applications Google offers (February 2012)

- ❏ App Engine
 - ❏ Run web applications in the Google infrastructure
 - ❏ SLA 99.95%
- ❏ Cloud Storage
 - ❏ Key-value storage
 - ❏ SLA 99.9%
- ❏ Prediction API
 - ❏ Pattern-matching and machine learning capabilities
 - ❏ SLA 99.9%
- ❏ ...

Google Cloud Services - App Engine I

Offers three environments where your application can run

☒ Java

- ☒ Java 6 JVM in safe sandboxed environment
- ☒ You can use standard JavaServer Pages (JSP)
- ☒ Some limits what application can do exist. You cannot:
 - ☒ Write to the file system
 - ☒ Open a socket or access another host directly
 - ☒ Create processes or threads directly or indirectly
 - ☒ Directly communicate with system via JNI
- ☒ There are limits for web application in general
 - ☒ Maximal size for web application and static files 150MB
 - ☒ Request and response size 32MB
- ☒ You can use also other languages compatible with JVM: Ruby, Scala, Jython, ...



☒ Python

- ☒ Python 2.5 or 2.7
- ☒ Similar limits like for Java environment (write to file system, use sockets, create threads, size of web application, ...)

Google Cloud Services - App Engine II

Supported environments (continue)

- ❏ Go
 - ❏ Go is a new Google's language
 - ❏ Similar limits like for Java or Python environment

Storage options

- ❏ App Engine DataStore
 - ❏ NoSQL database
 - ❏ Build on top of Google's BigTable (used by web search and Gmail)
 - ❏ Maximal size for item 1MB
- ❏ Cloud SQL
 - ❏ Based on MySQL
 - ❏ Maximal database size is 10GB
- ❏ Cloud Storage

Google Cloud Services - Cloud Storage

- ❏ Key-value storage for big objects (“bucked” 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
- ❏ You can specify location of bucket: US or Europe
- ❏ Data accessible as
 - ❏ <http://commondatastorage.googleapis.com/bucket/object>
 - ❏ <http://bucket.commondatastorage.googleapis.com/object>
- ❏ Redirects via DNS CNAME possible - <http://yourdomain/object>
- ❏ Support access control
- ❏ Supports resumable uploads
- ❏ Supports streaming

Google Cloud Services - Instance classes

Frontend instance

- ☒ Scales up and down automatically as number of requests increase / decrease

Backend instance

- ☒ Doesn't scale automatically - you specify number of instances you want to use

Instance	Memory limit	CPU limit
Frontend F1 (default)	128MB	600MHz
Frontend F2	256MB	1.2GHz
Frontend F4	512MB	2.4GHz
Backend B1	128MB	600MHz
Backend B2 (default)	256MB	1.2GHz
Backend B4	512MB	2.4GHz
Backend B8	1024MB	4.8GHz

Google Cloud Services - Pricing Instances

Prices (February 2012)

Instance	Cost per hour per instance
Frontend F1	\$0.08
Frontend F2	\$0.16
Frontend F4	\$0.32
Backend B1	\$0.08
Backend B2	\$0.16
Backend B4	\$0.32
Backend B8	\$0.64

- 28 on-demand frontend instance hours per day are for free
- 9 backend instance hours per day are for free

Google Cloud Services - Pricing Cloud Storage

Prices (February 2012)

Monthly usage	Storage price (per GB per month)	Network outbound (per GB per month) for USA, Europe	Network outbound (per GB per month) for Asia
First 0 - 1TB	\$0.13	\$0.12	\$0.21
Next 9TB	\$0.12	\$0.11	\$0.18
Next 90TB	\$0.105	\$0.08	\$0.15
Additional storage	Individual	Individual	Individual

- Inbound data are for free
- PUT, POST, list objects, list buckets \$0.01 per 1000 requests / month
- GET, HEAD requests \$0.01 per 10000 requests / month

Google Cloud Services - Free trial quota

First project that uses Google Cloud Storage will get for free (February 2012)

- ❏ 28 on-demand frontend instance hours per day
- ❏ 9 backend instance hours per day
- ❏ 5GB of storage
- ❏ 25GB download of data per month
- ❏ 30000 GET, HEAD requests per month
- ❏ 3000 PUT, POST, list objects, list buckets per month

Note: Users signing before 2011-05-10 were getting 100GB of free storage (valid till 2011-06-30).

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 scalability 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 and vendor
- ❏ 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
- ❏ For authors of cloud applications platforms provide less functionality than existing platforms
 - ❏ Requires new know-how for the whole team

Web catching failures of all major cloud platforms

- ❏ <http://cloudfail.net>

Děkuji za pozornost.

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



INVESTICE DO ROZVOJE VZDĚLÁVÁNÍ