

T-110.5121 Mobile Cloud Computing Public Infrastructure as a Service (IaaS) & Assignment 1 21.09.2011

Yrjö Raivio & Ramasivakarthik Mallavarapu
Aalto University, School of Science
Department of Computer Science and Engineering
Data Communications Software
Email: yrjo.raivio(at)aalto.fi
Course email: t-110.5121(at)tkk.fi



Cloud computing and mobile networks

NSN to push cloud computing to telco gear market

 Recommend  38 people recommend this. Be the first of your friends.

By **Tarmo Virki**, European Technology Correspondent
HELSINKI | Mon Sep 19, 2011 7:29pm EDT

(Reuters) - Nokia Siemens Networks (NSN) unveiled on Tuesday its new Liquid technology which promises to ease telecom network congestion and possibly change the market for such equipment for good.

The technology from the world's No. 2 mobile gear maker after Ericsson shares network resources across mobile and broadband networks, similar to cloud computing.

"For operators it offers a better use of capital. You invest in total network load, not in peak levels," Phil Twist, head of marketing at NSN's Network Systems unit, said in an interview.

Telecom operators around the world are struggling with a shortage of network capacity as the use of video on smartphones and tablets proliferates. The shortages usually occur during rush hours at specific locations.

Networks at city centers are crowded, base stations in suburbs which are scarcely used during the daytime. With traditional technology, operators have to buy new base stations for the centers to cope with the growth.

According to Nokia Siemens, up to 80 percent of base stations' processing capacity and up to half of core networks' capacity is unused.

CHANGING THE MARKET?

 Tweet 144

 Share 90

 Share this

 Email

 Print

Related News

Chipmaker
Broadcom to buy
NetLogic for \$3.7
billion
Mon, Sep 12 2011

Internap not for sale,
bets on data center
growth: CEO
Fri, Sep 9 2011

Baidu offers glimpse
of new mobile OS
Fri, Sep 2 2011

Insight: Jobs exit
opens door for
nimble Apple rivals
Thu, Aug 25 2011

Analysis: Network
gear makers
stumble, but growth
still there

Source: Reuters and <http://www.nokiasiemensnetworks.com/portfolio/liquidnet>



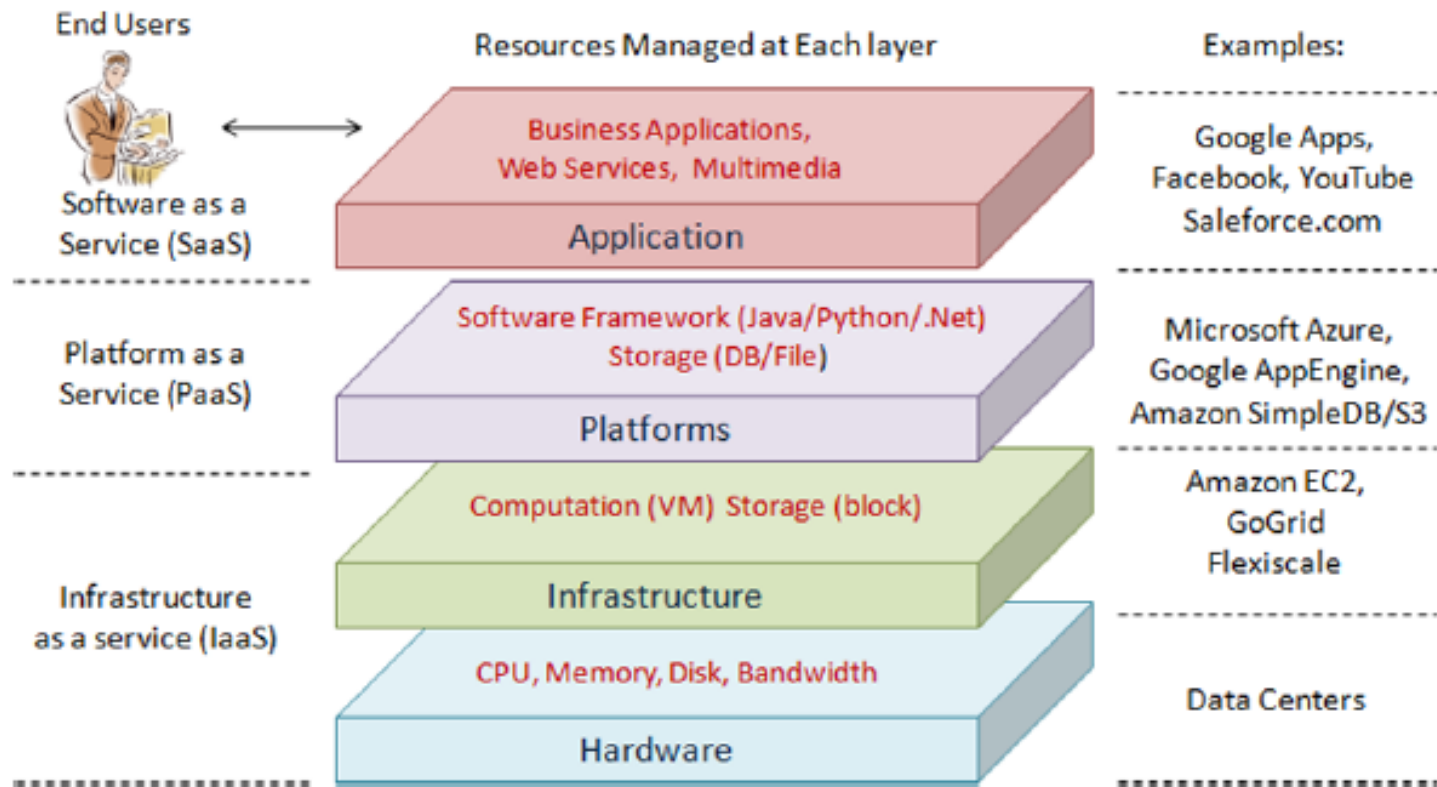
Outline

- **Public Infrastructure as a Service (IaaS)**
 - Types
 - Amazon EC2
 - Rackspace
 - RightScale
 - PaaS options
- **Amazon EC2 in practice**
- **Assignment 1**



Public Infrastructure as a Service (IaaS)

IaaS, PaaS, SaaS



Source: Qi Zhang, Lu Cheng and Raouf Boutaba, Cloud computing: state-of-the-art and research challenges, 2010.




IaaS cloud types

- **Private cloud**
 - Host own resources
 - Provide to internal customers only
 - Provision with cloud interfaces
- **Public cloud**
 - Resources made available to the general public via the internet
 - Scalable
 - Pay for what you use
- **Hybrid cloud**
 - Some resources provided internally and some resources outsourced
- **Community cloud**
 - Address special needs of a community



Amazon Elastic Compute Cloud (EC2) growth has been exponential



Source: RightScale Blog, available at <http://blog.rightscale.com/2009/10/05/amazon-usage-estimates/>



Amazon EC2

- **Various infrastructure Web Services. Most famous**
 - Amazon Elastic Compute Cloud (Amazon EC2)
 - Various operating systems, databases, applications
 - Amazon Simple Storage Service (Amazon S3)
- **Other:**
 - CloudFront (content delivery for the users for the nearest location to them)
 - Relational Database Service (RDS) (MySQL or Oracle)
 - SimpleDB
 - Simple Notification Service (SNS) e.g e-mail
 - Simple Queue Service (SQS) (application message exchange)
 - Virtual Private Cloud (VPC): to connect existing resources to AWS resources via VPN including IP ranges and routes, NAT, filtering



Amazon EC2 Pros and Cons

- **Pros**

- No lock-in. Relatively easy to move code to another box.
- Code in C#, .NET, ASP.NET MVC / Visual Studio
- SQL Server (Express/Compact)
- Amazon staff very active in the community

- **Cons**

- No free quota - minimum cost per month, even if 0 visitors
- Scaling harder (if required) - need to handle data spread across n instances
- Single point of failure with single Micro instance

Source: <http://blog.dantup.com/2010/10/google-app-engine-gae-vs-amazon-elastic-computing-ec2-vs-microsoft-azure>



Amazon EC2 price components

- **Zones**
- **Instance size**
- **Storage size**
- **Reserved instances**
- **Spot instances**
- **Data transfer**
- **Elastic IP address**
- **Monitoring services**
- **Elastic load balancing**
- **VPN**

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



Rackspace

- **Competitor for Amazon**
- **Cloud Servers**
 - Memory 256 MB.. 15,5 GB / host
 - Linux or Windows Server 2008 / 2003
 - Load balancer
 - Starting \$0.015 / h
- **Cloud Files, built on OpenStack**
 - Online storage from \$0.15 / GB
 - Each server adds 10 GB to the storage pool
 - Bandwidth out \$0.18 and in \$0.08 (GB)

Feature comparison

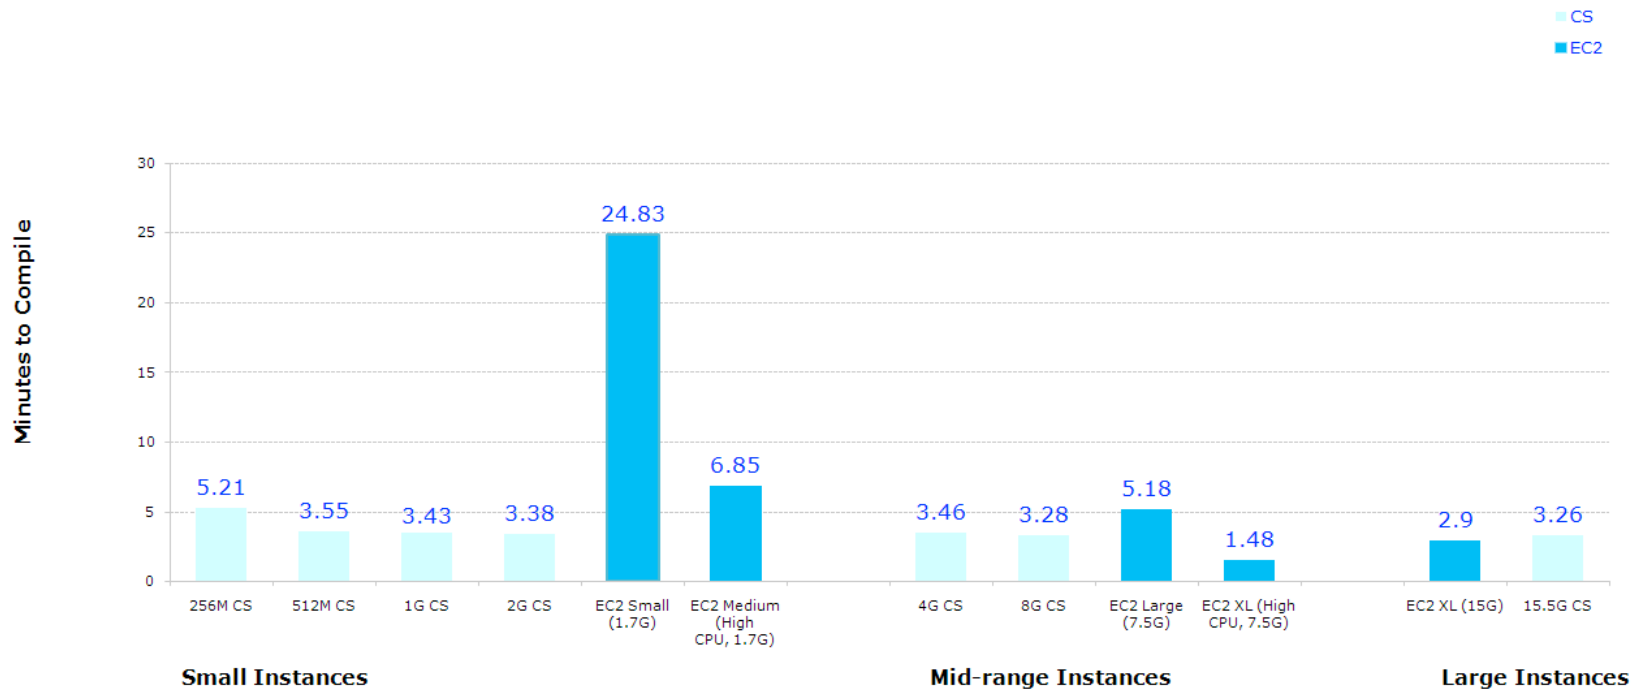
#	Feature	Rackspace	Amazon
1	Support	Included	400\$/month or 10%
2	3rd party SW	Linux, MS, MySQL..	No
3	Admin help	Yes	No
4	Persistence	Yes	No/EBS
5	Server sizes	All sizes	No low end support
6	Hybrid hosting	Cloud&Hosting	Only VPN
7	CPU scheduling	Free if available	Capped CPU
8	Compute power	2 x better than EC2	Cheaper instance
9	Disk I/O	Better than EC2	EBS improves
10	IP addresses	No NAT	NATed
11	Openness	OpenStack, DMTF	Not active

Source: http://www.rackspace.com/cloud/cloud_hosting_products/servers/compare/



Performance comparison

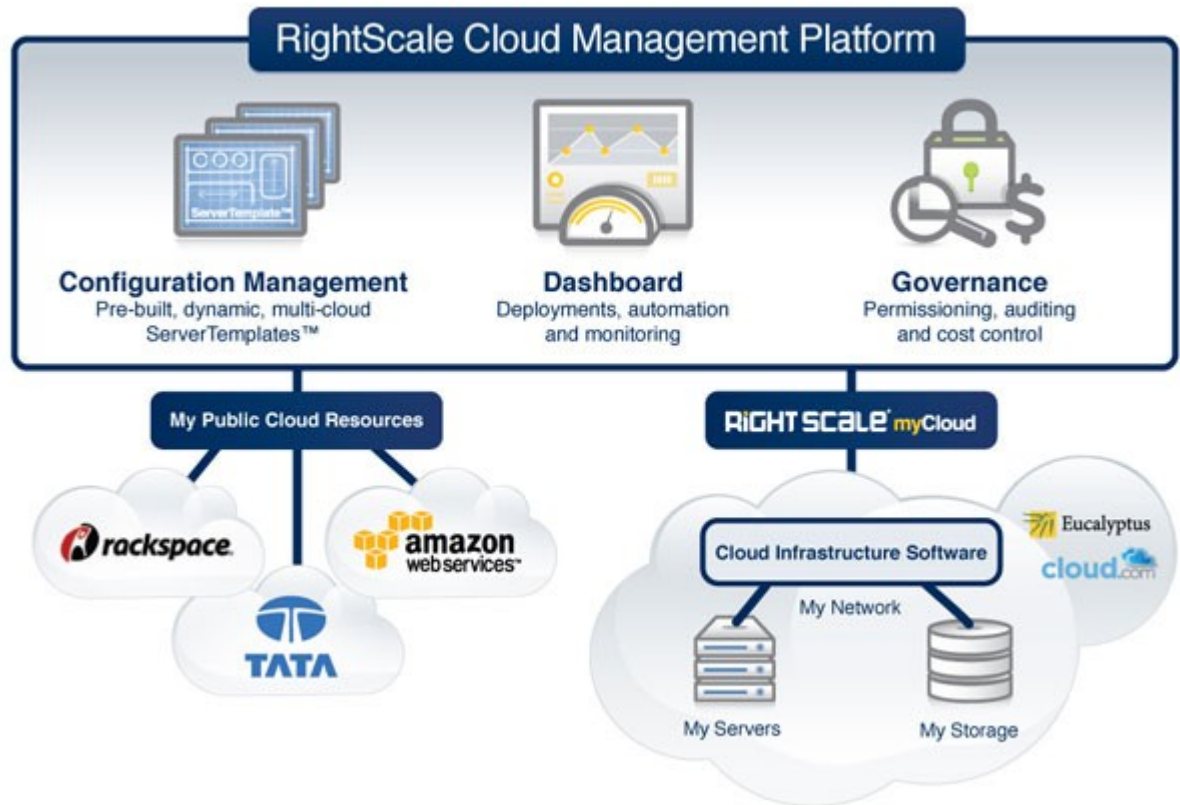
CloudServers vs. Amazon EC2 Linux Kernel Compilation Time Results



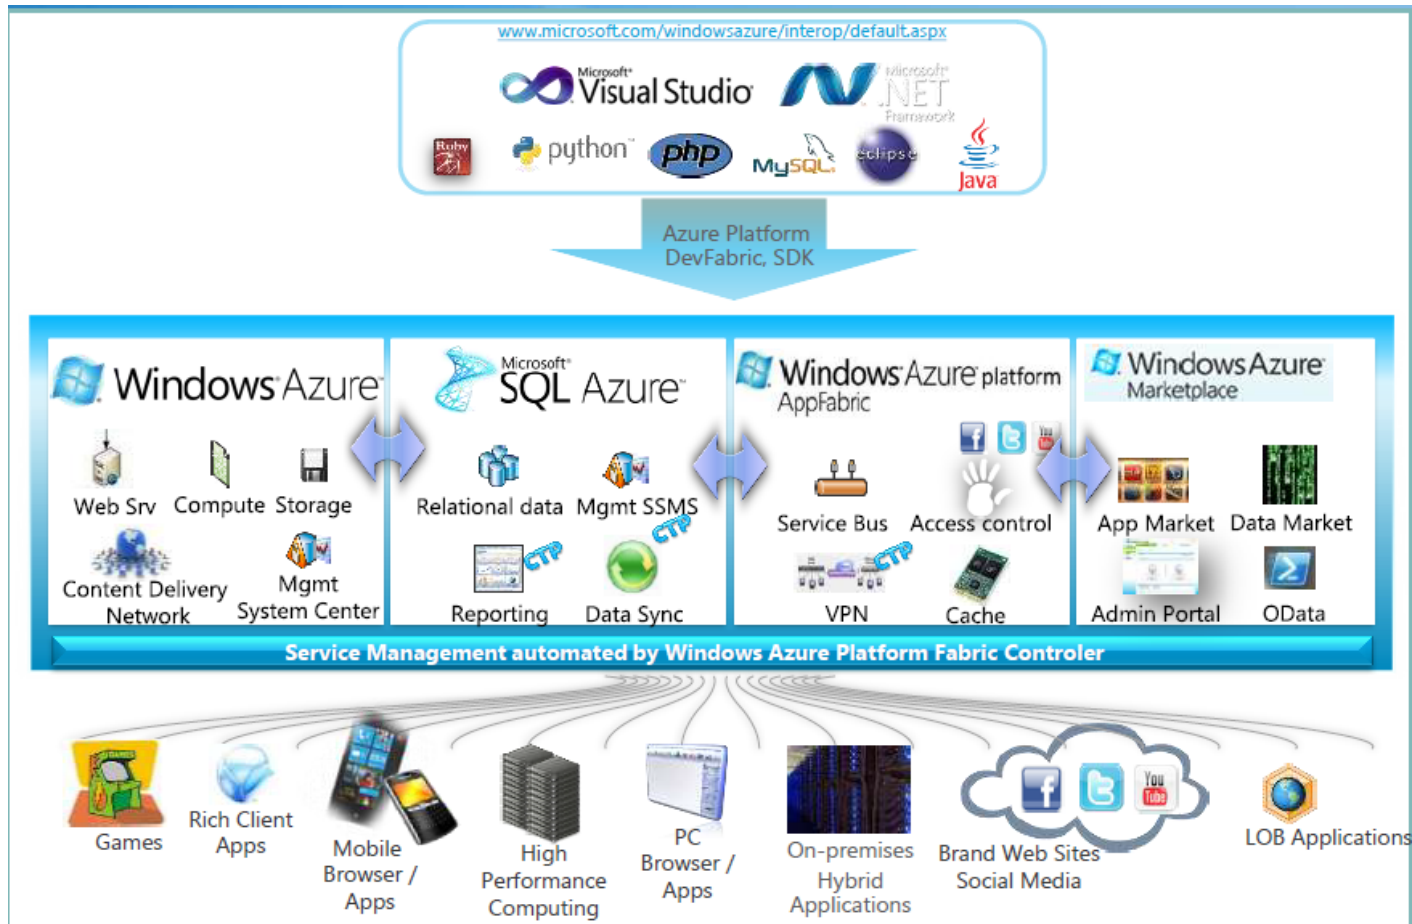
Source: <http://www.thebitsource.com/featured-posts/rackspace-cloud-servers-versus-amazon-ec2-performance-analysis/>

RightScale

- Intermediate layer between customer and cloud provider
- Hybrid cloud
- Generic management interface to all cloud resources



Microsoft Azure



Source: Pasi Mäkinen, Cloud Evangelist, Microsoft, EIT Cloud Computing Summer School June 2011



Components

	Windows Azure PaaS
Applications	Windows Azure Service Model
Runtimes	.NET 3.5/4, ASP .NET, PHP
Operating System	Windows Server 2008-Compatible OS
Virtualization	Windows Azure Hypervisor
Server	Microsoft Blades
Database	SQL Azure
Storage	Windows Azure Storage (Blob, Queue, Table)
Networking	Windows Azure-Configured Networking

Pricing

X-Small (XS) \$0.05 usage hour	Small (S) \$0.12 usage hour	Medium (M) \$0.24 usage hour	Large (L) \$0.48 usage hour	X-Large (XL) \$0.96 usage hour
Extra Small 1 x 1.0GHz 5Mbps IO 768 MB RAM 20 GB disk (transient storage)	Small 1 x 1.6GHz 100Mbps IO 1.75 GB RAM 225 GB disk (transient storage)	Medium 2 x 1.6GHz 200Mbps IO 3.5 GB RAM 500 GB disk (transient storage)	Large 4 x 1.6GHz 400Mbps IO 7.0 GB RAM 1000 GB disk (transient storage)	X-Large 8 x 1.6GHz 800Mbps IO 14.0 GB RAM 2000 GB disk (transient storage)

Billing starts when application is deployed and stops when it's stopped AND deleted!



Own Azure account

- **Building VM Role Step-by-step:**
 - <http://msdn.microsoft.com/en-us/gg502178>
- **Free 30 day account:**
 - <https://windowsazurepass.com/>
 - Code UMBFIN
- **Longer term trial:**
 - <http://www.microsoft.com/windowsazure/free-trial/>
 - Does require credit card for usage above free quota
- **Paid accounts:**
 - <http://www.microsoft.com/windowsazure/offers/buy.aspx>

Source: Pasi Mäkinen, Cloud Evangelist, Microsoft, EIT Cloud Computing Summer School June 2011



Google Apps Engine

- **Pros**

- Generous free quota
- Billed for actual CPU usage, not "live" hours
- Scaling is (relatively) easy
- "Free" built-in Google authentication

- **Cons**

- (Pretty much) locked in to Google App Engine See comments below
- No support for C#/.NET
- Non-relational datastore - more restrictive on how data can be accessed
- (Relatively) high rate of datastore errors, need to code to handle them
- Blog is already hosted here - don't need to do any work ;)
- No local storage, so unable to use libraries that rely on writing to disk

Source: <http://blog.dantup.com/2010/10/google-app-engine-gae-vs-amazon-elastic-computing-ec2-vs-microsoft-azure>



Force.com

- **220 000 applications, 97 000 companies**
- **99.9% availability**
- **Integrated sandbox, APIs, visual tools**
- **Social network: components, sharing**
- **Mobile support, any device**



”Decision algorithm”

If (development_environment = .NET) Then

CloudHost = Azure

If (server_requirement = Website) Then

If (new_website_build) Then

CloudHost = AppEngine

Elsif (money = Abundant) Then

CloudHost = Rightscale_EC2

Else

CloudHost = Rackspace

If (server_availability = 24x7) And (server_load = Constant) Then

CloudHost = Rackspace

Else

CloudHost = EC2

Source: <http://www.distractable.net/tech/amazon-aws-ec2-vs-rackspace-high-level-comparison/>



Additional reading material

1. EIT Cloud Computing Summer School, 6.-10.6.2011, available at: <http://cse.aalto.fi/studies/eit-2011/lecture-materials/>
2. Qi Zhang, Lu Cheng and Raouf Boutaba, Cloud computing: state-of-the-art and research challenges, J Internet Serv Appl (2010) 1: 7–18.
3. Alexander Lenk, Markus Klems, Jens Nimis, Stefan Tai and Thomas Sandholm , What's Inside the Cloud? An Architectural Map of the Cloud Landscape, CLOUD'09, May 23, 2009, Vancouver, Canada.
4. Neal Leavitt, Is Cloud Computing Really Ready for Prime Time? Computer, January 2009.
5. Ang Li, Xiaowei Yang, Srikanth Kandula and Ming Zhang, CloudCmp: Comparing Public Cloud Providers, IMC'10, November 1–3, 2010, Melbourne, Australia.
6. Vladimir Stantchev, Performance Evaluation of Cloud Computing Offerings, ADVCOMP 2009, October 11-16, 2009 - Sliema, Malta.

Amazon EC2 in practise Assignment 1

Course Registration System deployed on Amazon EC2

- Task:
 - Get some input through a HTML form
 - Query the MySQL database
 - Perform some trivial processing and
 - Finally display the result
- Initial Setup:
 - Install all the required software such as MySQL, Apache, JDK or PHP
 - Setup tables in the database (MySQL scripts will be provided)
- Input:
 - Application takes course codes as input
- Output:
 - Display course information like lecture and examination schedule



Additional information

- Problem statement:
 - Given course codes as input, query the database for lecture dates, examination dates and display the information
 - Output must also include if there is a clash in course schedules (overlapping of lecture times or examination times..please report!)
- Allowed programming languages:
 - PHP,Java
- HTML template and MySQL insert scripts will be provided to the students
- Some references for the programming part will also be posted in Noppa



Learning goals for Assignment 1

- Assignment 1 is **not** about:
 - Learning how to program in Java or PHP
 - Learning about databases
 - HTML formatting etc.
- Goals:
 - Practical experience on Amazon EC2
 - Amazon machine images
 - Amazon security groups and elastic IPs
 - Instance types, key pairs etc.



Questions?

Contacts:

Teacher: yrjo.raivio(at)aalto.fi, A122

Assistants: ramasivakarathik.mallavarapu(at)aalto.fi, PlayRoom

koushik.annapureddy(at)aalto.fi, A118

Course staff: t-110.5121(at)tkk.fi