

Cloud/Grid Computing


- Provide an introduction to cluster, grid and cloud infrastructures.
- Define an example of grid computing, and its advantages.
- Show an example of using a Cloud Infrastructure.
- Define the usage of Amazon Web Services, including S3 and the EC2 Cloud.
- Outline the usage of RESTful Web Services.




Cloud Computing




Introduction





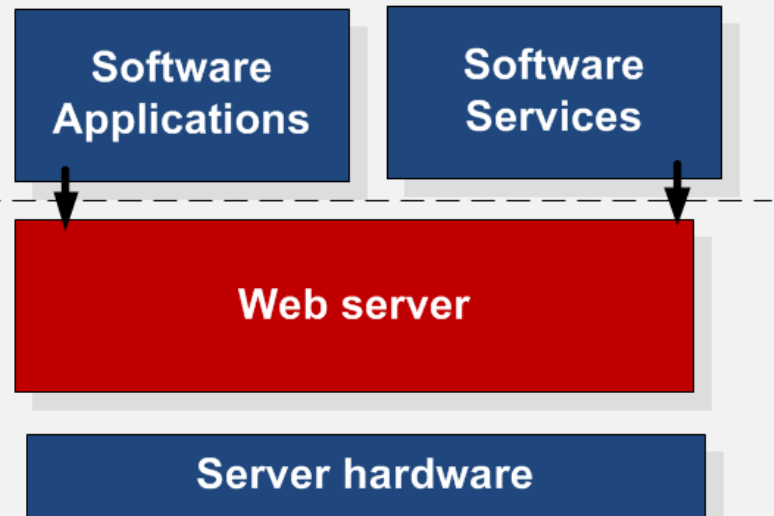
Current (Thick client where applications run locally)

Isolation between images




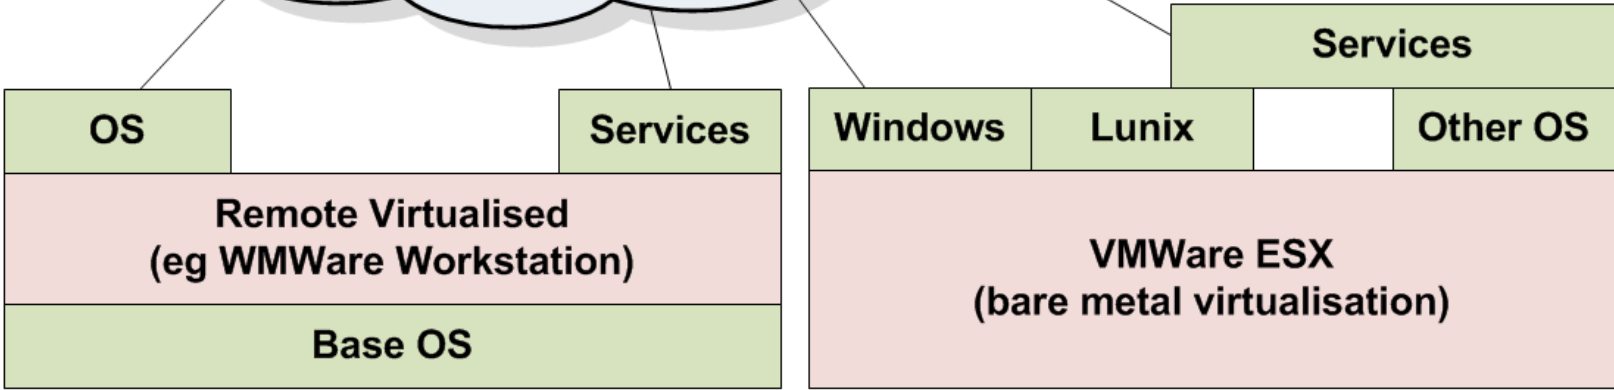

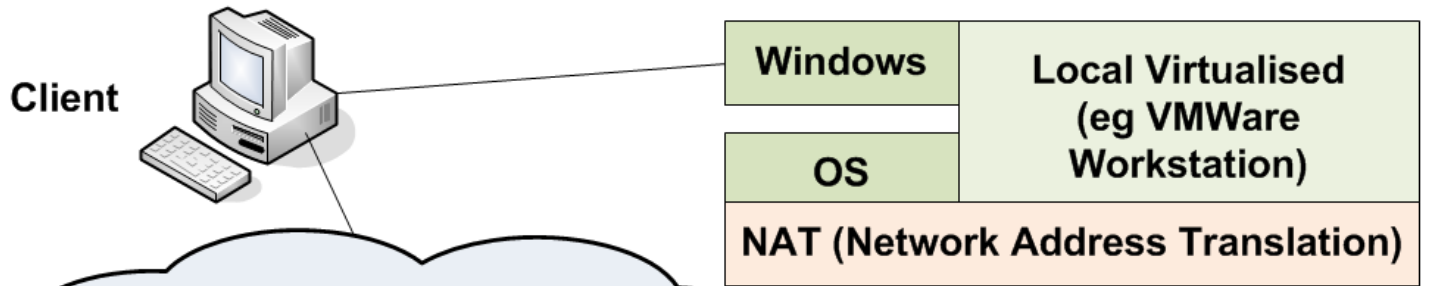
Virtualisation images allow multiple OSs to run on a single OS.

Future (Thin clients and Virtualisation)

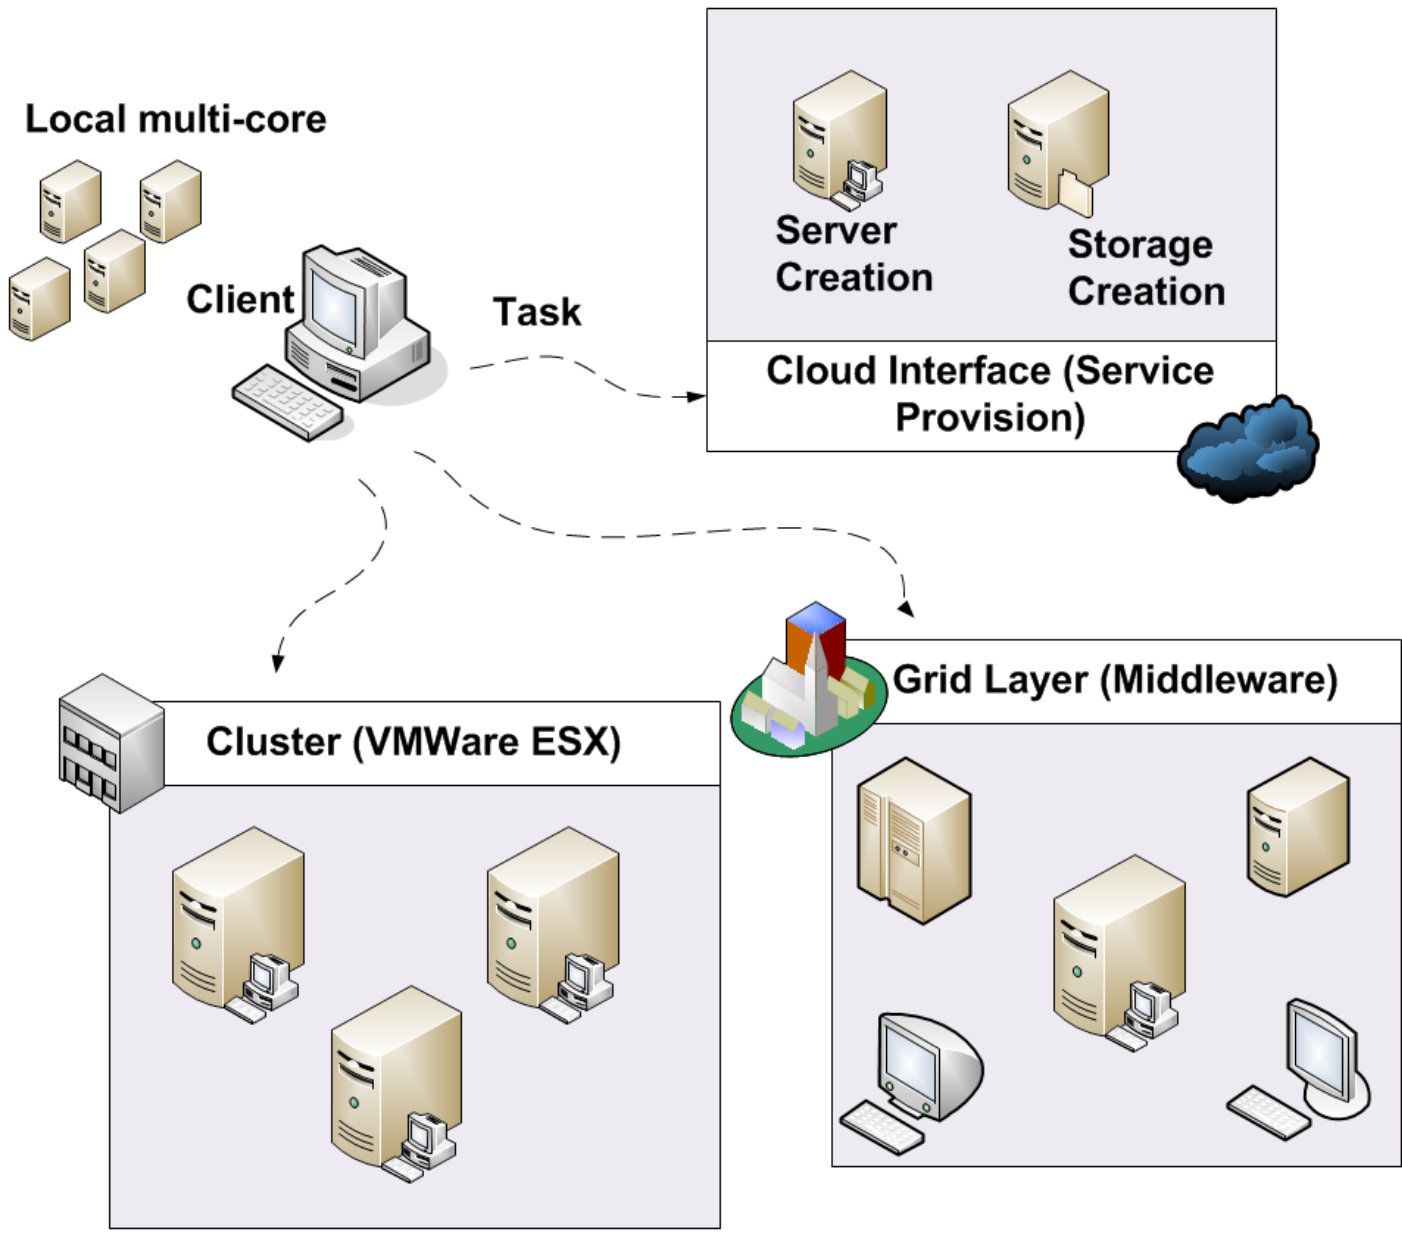


Virtualisation

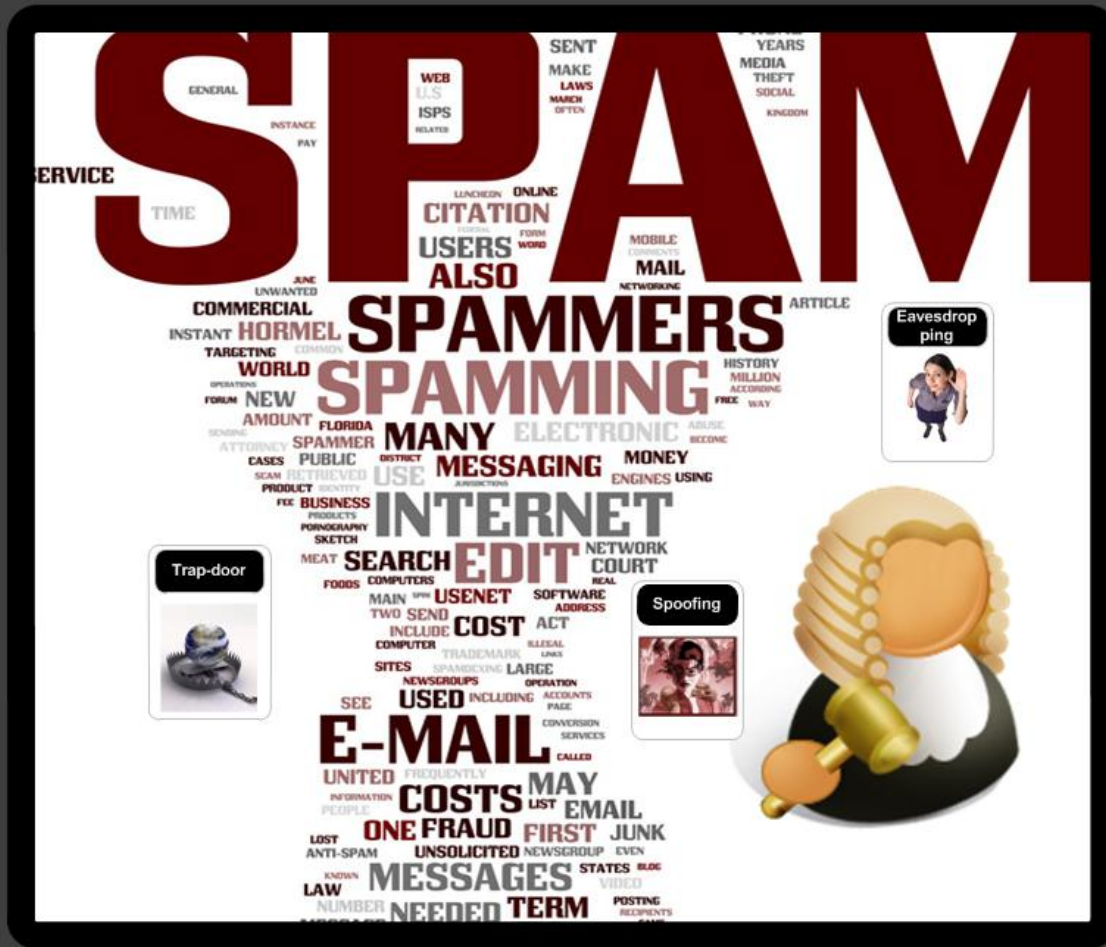
Cloud



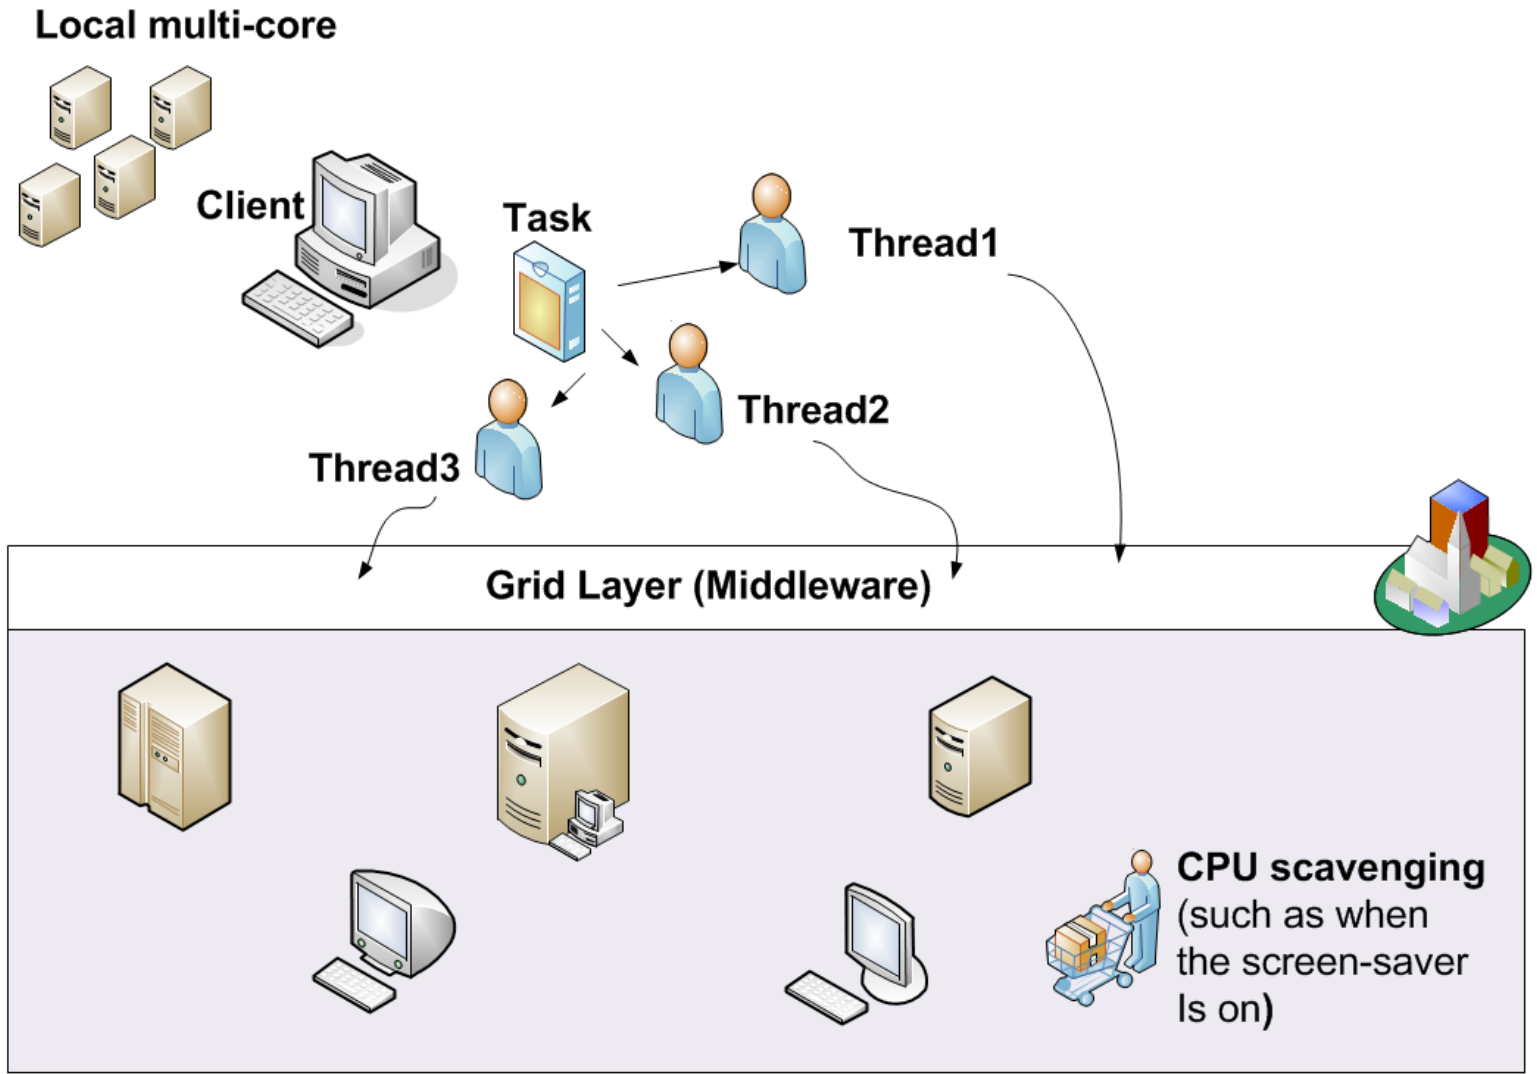
Virtualisation



Cloud Computing



Grid Computing




Client



1. Submit
a job





Data Manager




5. Read results



2. Push data set

Messaging
QueueWorkers watch
queue4. Publish
result3. Pull data
set

CPU scavenging
(such as when
the screen-saver
is on)





Okay... we select a **64-bit key** ...
which has 1.84×10^{19}
combinations

18.4 million million million different keys
000000000000....00000000000000000000
To
111111111111....11111111111111111111


How long will it take to cracked It by brute-force
(on average)?



A 64-bit key has 1.84×10^{19} combinations and it could be cracked by brute-force in 0.9×10^{19} goes.

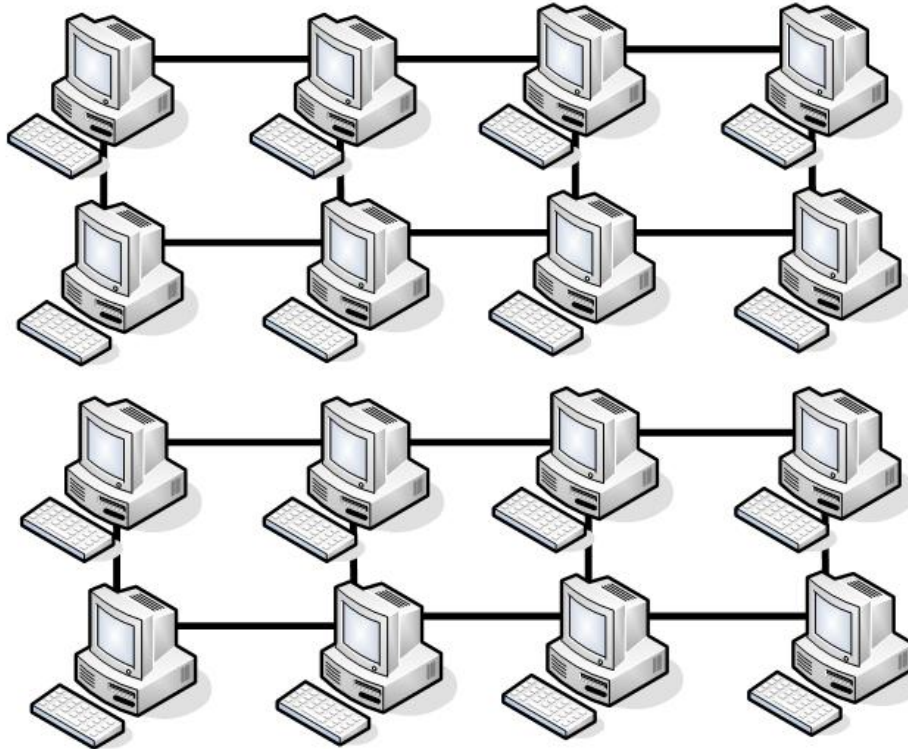


If we use a fast computer such as 1GHz clock (1ns), and say it takes one clock cycle to test a code, the time to crack the code will be:

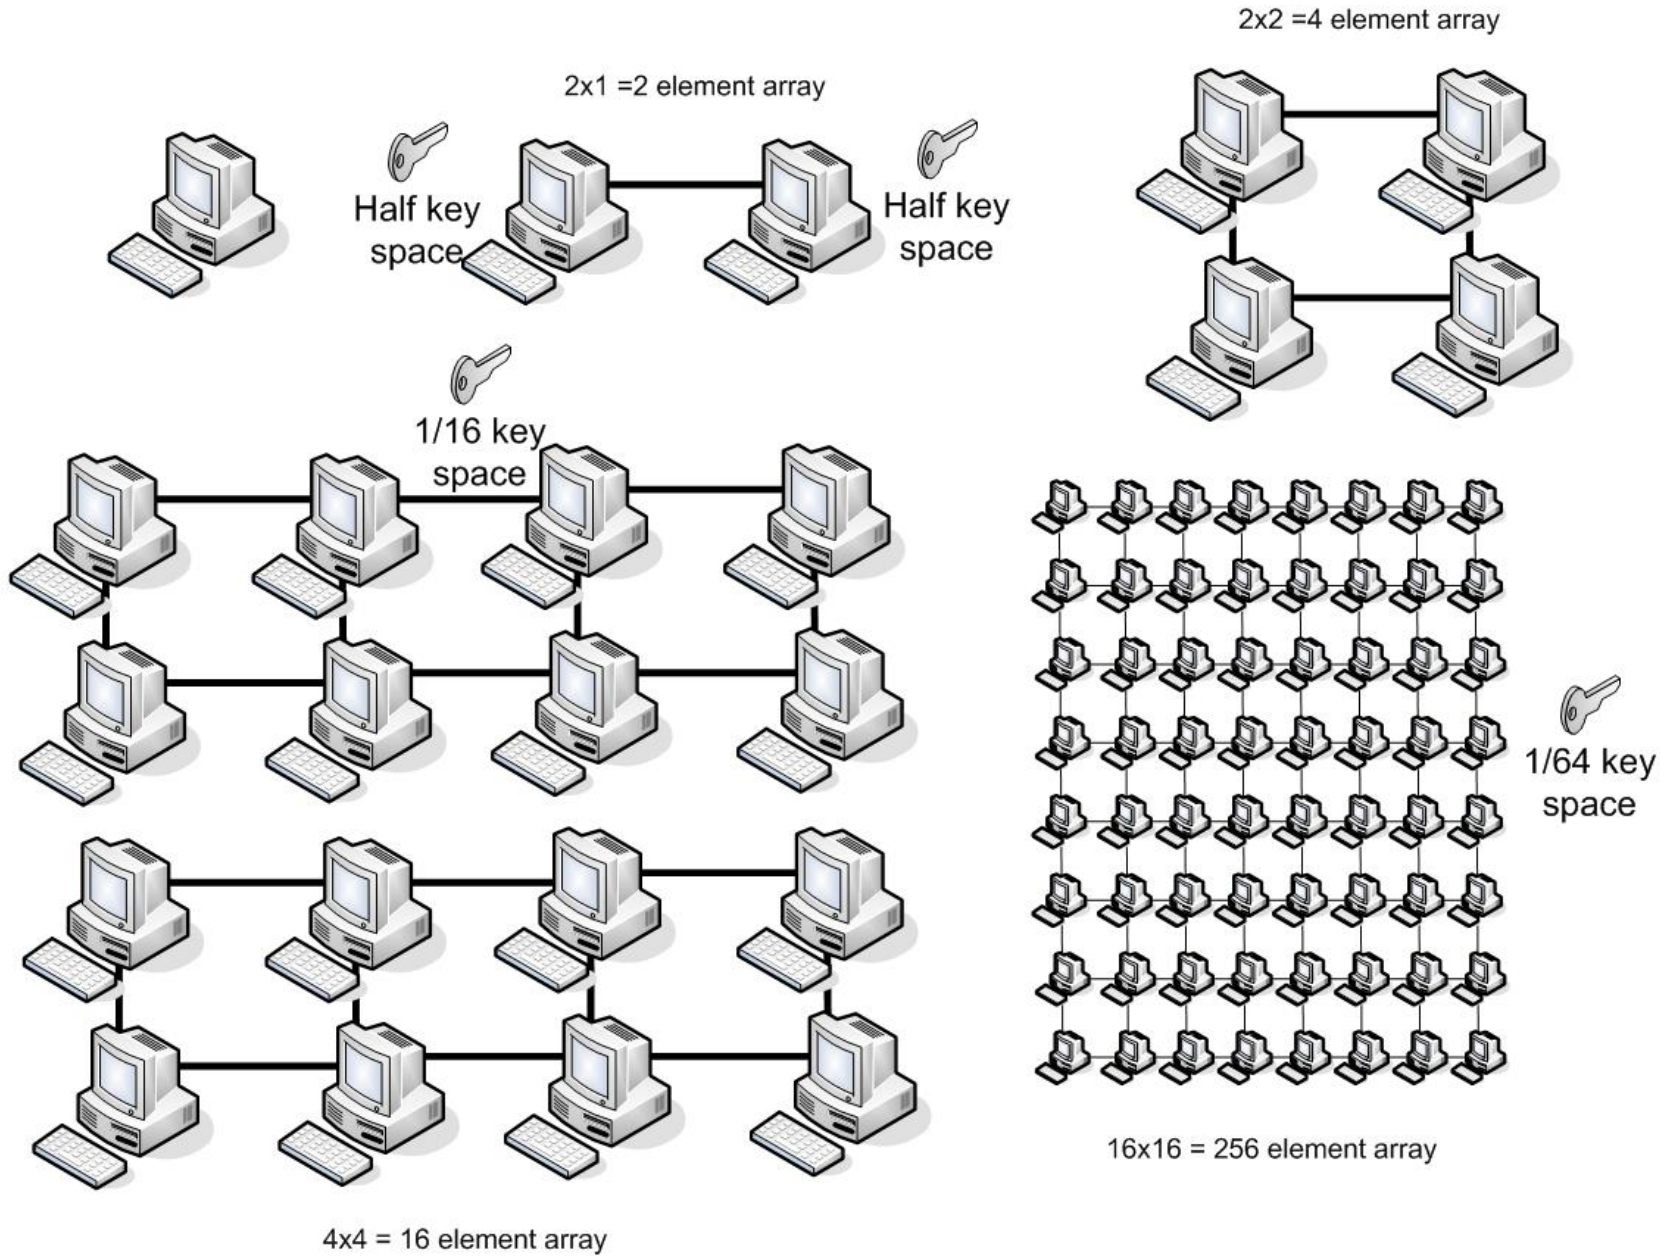
9,000,000,000 seconds (150 million minutes)
... **2.5 million hours** (285 years)



Date	Hours	Days	Years
2008	2,500,000	104,167	285
2009	1,250,000	52,083	143
2010	625,000	26,042	71
2011	312,500	13,021	36
2012	156,250	6,510	18
2013	78,125	3,255	9
2014	39,063	1,628	4
2015	19,532	814	2
+8	9,766	407	1
+9	4,883	203	1
+10	2,442	102	0.3
+11	1,221	51	0.1
+12	611	25	0.1
+13	306	13	0
+14	153	6	0
+15	77	3	0
+16	39	2	0
+17	20	1	0

56-bit DES:
 Developed
 1975
 30 years ago!
 ... now easily
 crackable

$2 \times 1 = 2$ element array $2 \times 2 = 4$ element array $4 \times 4 = 16$ element array $16 \times 16 = 256$ element array




2x2 =4 element array


2x1 =2 element array



Half key space



Half key space



Processors	Year 0	Year 1	Year 2	Year 3	Year 4	Year 5
1	104000 days	52000	26000	13000	6500	3250
4	26000	13000	6500	3250	1625	813
16	6500	3250	1625	813	407	204
64	1625	813	407	204	102	51
256	406	203	102	51	26	13
1024	102	51	26	13	7	4
4096	25	13	7	4	2	1
16,384	152hr	76hr	38hr	19hr	10hr	5hr
65,536	38hr	19hr	10hr	5hr	3hr	2hr
262,144	10hr	5hr	3hr	2hr	1hr	
1,048,576	2hr	1hr				

key space

16x16 = 256 element array

4x4 = 16 element array

Author: Prof Bill Buchanan

Grid

Cloud



1997. RSA Lab's 56-bit RC5 Encryption Challenge
 - 250 days and 47% of the key space tested) –
distributed.net



1998. RSA Lab's 56-bit DES II-1 Encryption Challenge - 39 days.
1998. RSA Lab's 56-bit DES II-2 Encryption Challenge - 2.5 days.

1999. RSA Lab's 56-bit DES-III Encryption Challenge - after 22.5 hours using EFF's Deep Crack custom DES cracker.



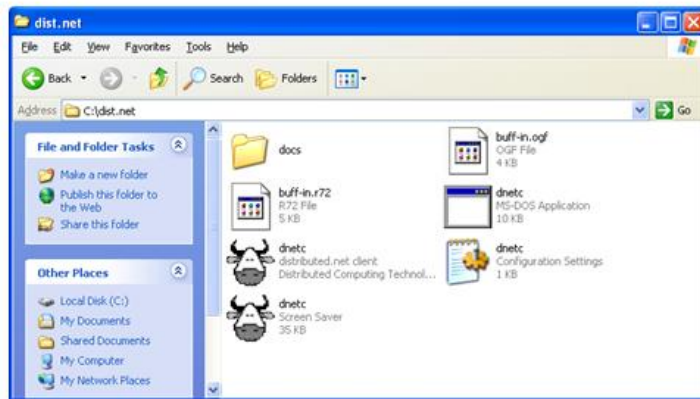
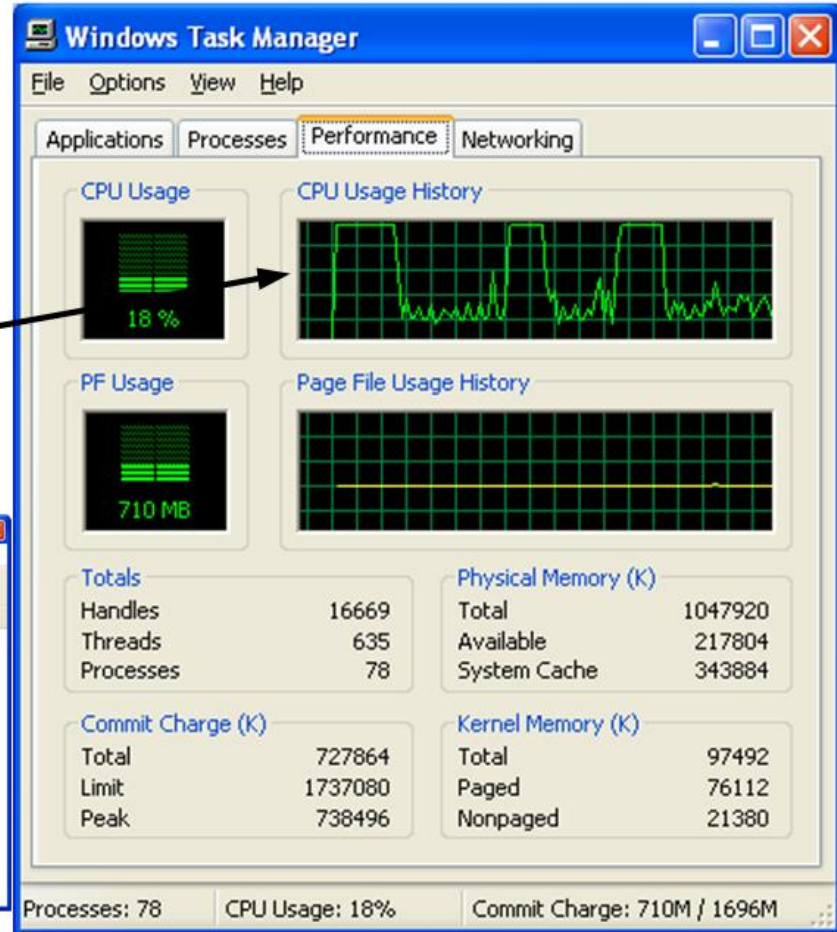
2002. RSA Lab's 64-bit RC5 Encryption Challenge
 — Completed 14 July 2002 – 1,757 days and 83%
 of the key space tested.

RSA Lab's 72-bit RC5 Encryption Challenge
 - In progress.







Distributed.net is starting and stopping (Max CPU when searching for possible keys)



Author: Prof Bill Buchanan




On-demand self-service. Consumers get server CPU, memory, bandwidth and storage resources whenever required.




Location independent resource pooling. Multiple customers use shared resources within the provider, without actually knowing where the exact location of these are.


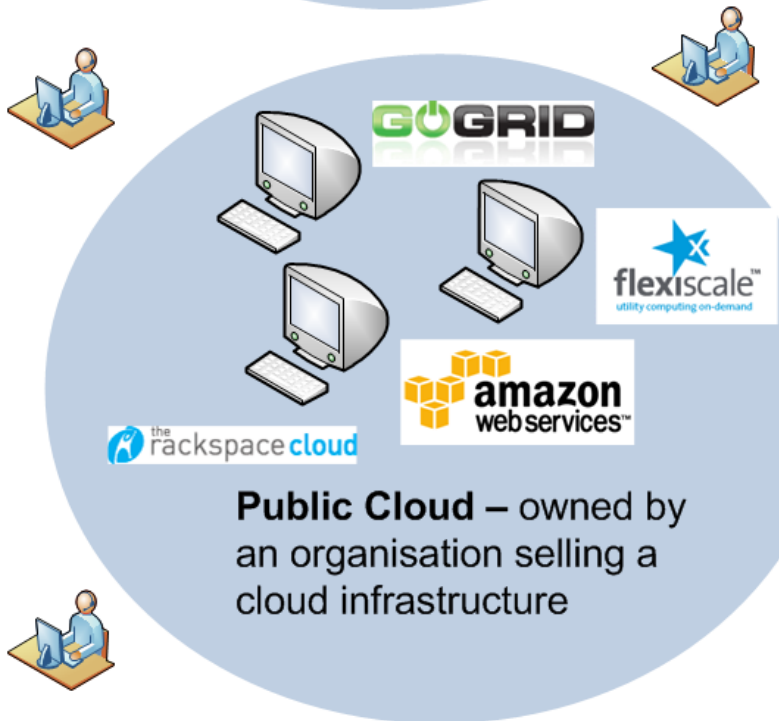


CLOUD





Rapid elasticity. Consumers can easily scale-up and scale-down, whenever required.



Pay per use. All access to resources is monitored, and paid for either by advertising or usage. Payment methods: per users created, per hour usage (service), etc.






External Cloud



Q: Rate the challenges/issues ascribed to the 'cloud'/on-demand model
 (1=not significant, 5=very significant)



Source: IDC Enterprise Panel, August 2008 n=244

Audit/compliance

Can I be compliant with statutory and regulatory requirements?

- Where is my data stored?
- Who handles breach notifications?
- How long is my data stored for?
- How is eDiscovery handled?



Client

Software as a Service (SaaS)

- User interface.
- Machine interface

Platform as a Service (PaaS)


- Service Oriented Architecture (SOA)
- Sophisticated Web Services
- Developing
- Testing
- Deploying
- Hosting
- Service platform providers, e.g. Google GAE, Microsoft Windows Azure

Infrastructure as a Service (IaaS)

- Resource virtualisation
- Computing power
- Storage capacity
- Network bandwidth
- Usage-based payment scheme
- Cloud enablers, e.g. Amazon EC2 / S3

Hardware as a Service (HaaS)

- Cluster & data centre providers
- Reduction of capital & operation investments
- Enhanced reliability – redundancy, replication & failover
- Enhanced scalability
- Enhanced load-balancing





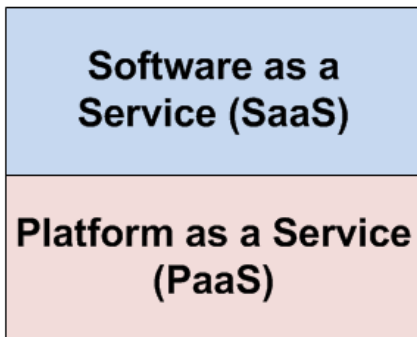
Servers

Cloud

Cloud



**API interface
To virtual servers**



Amazon Web Services

Cloud

Cloud

IaaS

Client



Pay-per-usage




API interface
To services

Microsoft Azure,
Google Application
Engine (GAE),
Amazon EC2

Software as a
Service (SaaS)

Platform as a Service (PaaS)

Infrastructure as a
Server
(IaaS)



SOAP, WSDL
Used to find services

Cloud

Cloud

PaaS

Client



Pay-per-usage

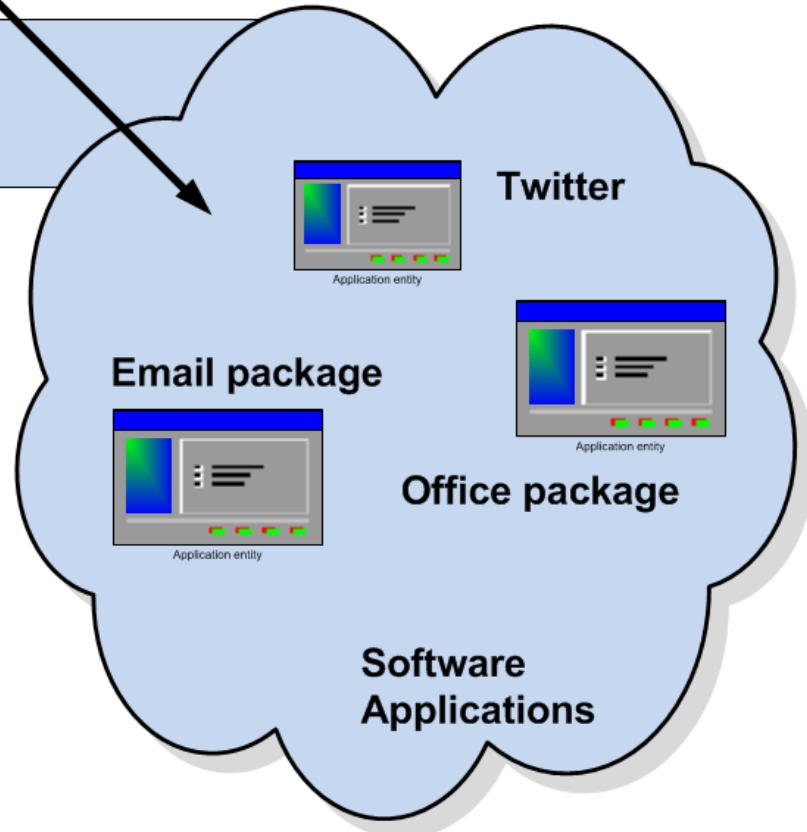


Access to applications

Software as a Service (SaaS)

Platform as a Service (PaaS)

Infrastructure as a Server (IaaS)



Cloud

Cloud

SaaS

Cloud Computing



Amazon EC2

Client



Amazon SimpleDB

This produces a mixture of structured data storage with the reliability of a traditional database.



Pay-per-usage

Amazon CloudFront

This allows content to be placed close to the places where it is to be consumed, the content thus gets moved to the edge of the cloud to support rapid delivery of content.

Amazon Elastic Cloud Compute (Amazon EC2)

This is the core of the Amazon Cloud, and provides a Web services API to create, manage and delete virtual servers within the Amazon Cloud. This includes US and European data centres, and uses the Xen hypervisor for the management of the servers.

Amazon Simple Storage Service (Amazon S3).


This provides data storage with web services through APIs. It differs from normal filesystems in that it does not have a hierarchical structure. Instead it uses buckets, which are unique namespaces across all of the Amazon customers. It is thus not a filesystem, and is a Web service, thus applications need to be written which specifically store data into the S3 Cloud.

Amazon Simple Queue Service (Amazon SQS).

This supports a grid infrastructure, where message can be passed to a queue, and then consumed by any subscribers.

AWS

Cloud



GET /resource/myservice

POST /resource/myservice

GET /resource/myservice/elem111

PUT /resource/myservice/elem111

DELETE /resource/myservice/elem111

List members in a collection

CREATE a new member in a collection

READ a member in a collection

UPDATE a member in a collection

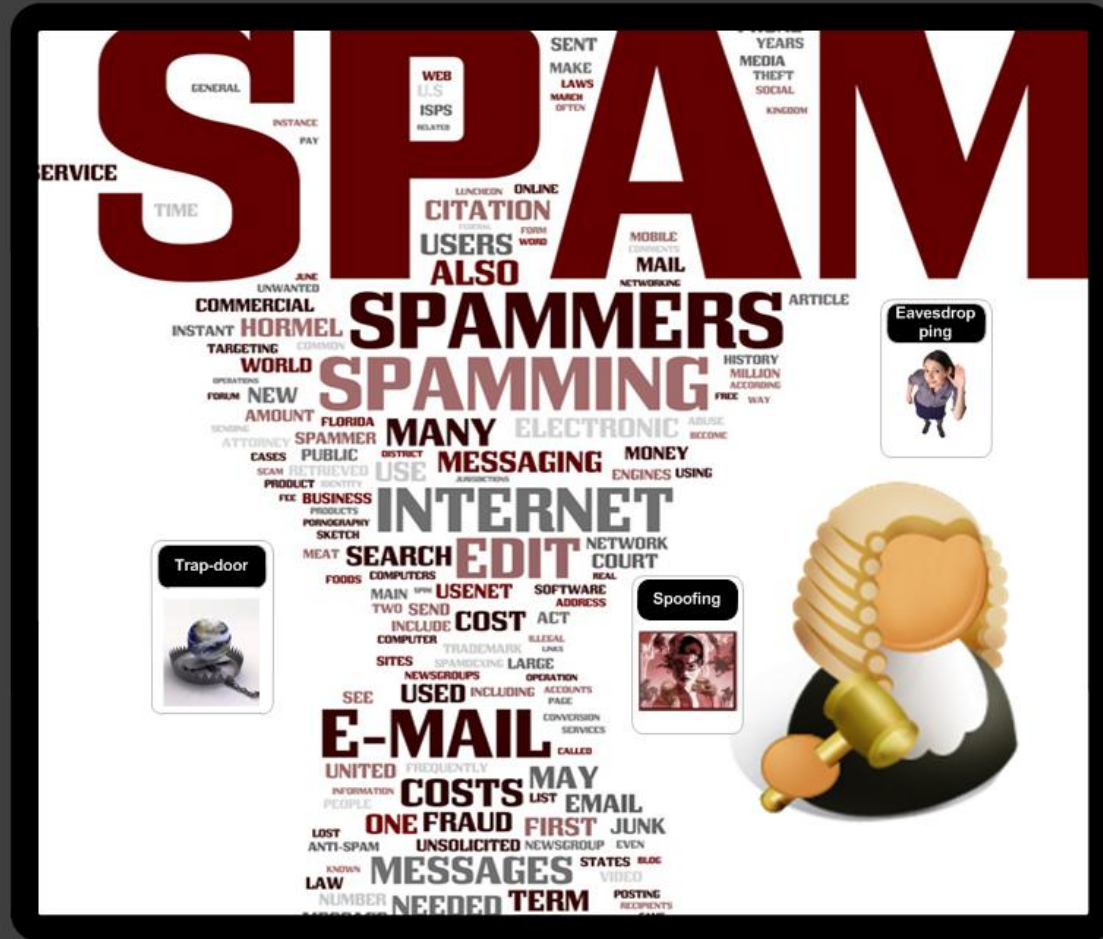
DELETE a member in a collection

- REST defines a set of architectural principles that focus on a system's resources.
- Is stateless, and focuses on how to define states using HTTP.
- Used by Yahoo, Facebook, Amazon, Microsoft Azure, etc.
- A lightweight alternative to SOAP/ WSDL.

Web Service

Create
Read
Update
Delete

Cloud Computing



Amazon E3

Client



amazon
web services™

Pay-per-usage

Amazon Simple Storage Service (Amazon S3).

This provides data storage with web services through APIs. It differs from normal filesystems in that it does not have a hierarchical structure. Instead it uses buckets, which are unique namespaces across all of the Amazon customers. It is thus not a filesystem, and is a Web service, thus applications need to be written which specifically store data into the S3 Cloud.

\$0.150 per GB - first 50 TB / month of storage used

```
s3cmd mb s3://bill.bucket
s3cmd put myfile.mp3 s3://bill.bucket/myfile.mp3
s3cmd get s3://bill.bucket/myfile.mp3 myfile.mp3

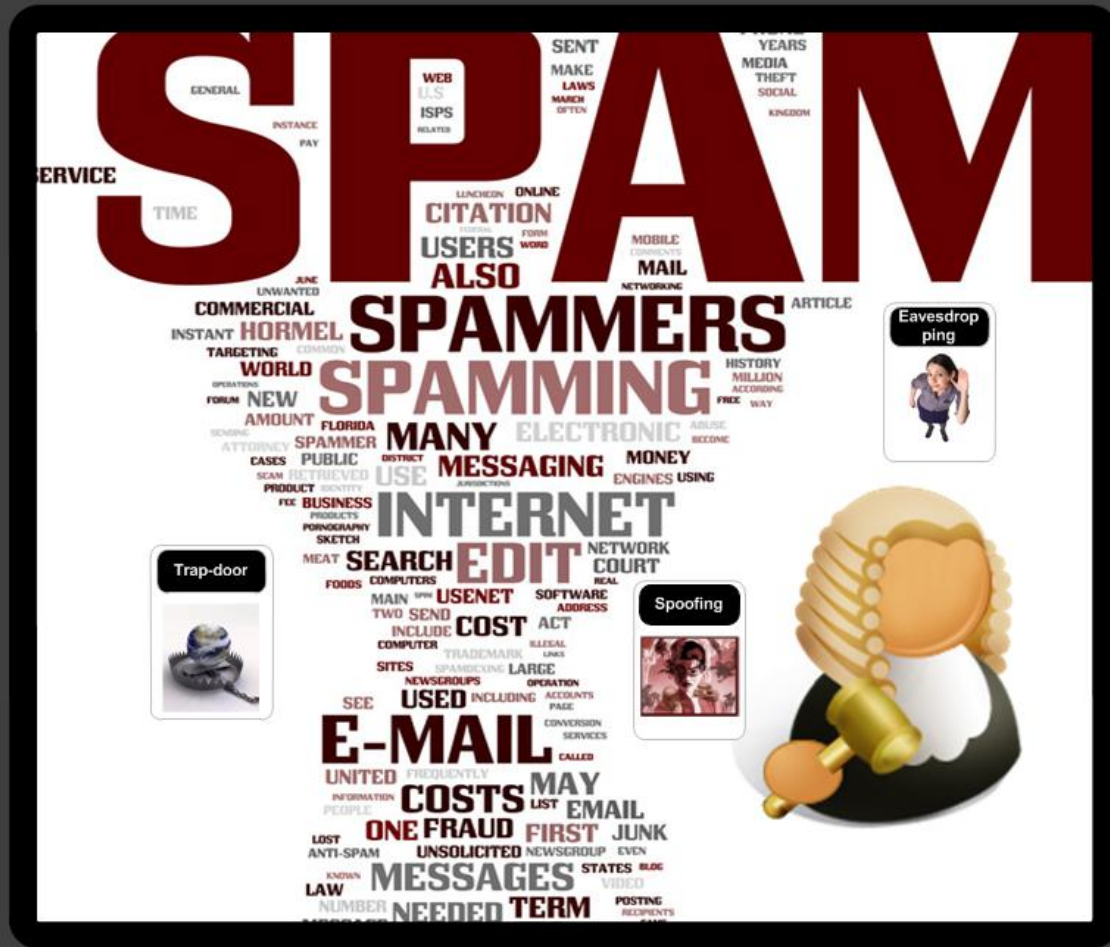
s3cmd ls
s3cmd ls s3://bill.bucket/

s3cmd delete s3://bill.bucket/myfile.mp3
s3cmd mb s3://bill.bucket
```

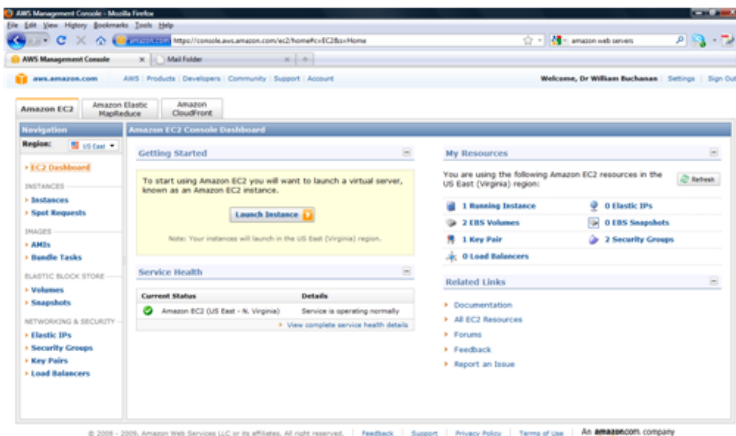
AWS

Cloud


Cloud Computing



Amazon EC2




Amazon Machine Images (AMI)



i-0d895567

Machine Instances

- Command line tools
- Web Service Console



i-0d895566

```
$ ec2-describe-images -o AKIAIWUMTTAZYST2I2AA
$ ec2-describe-images
IMAGE ami-45c22e2c powerdns/image.manifest.xml 495219933132 available private
$ ec2-run-instances i-0d895566
```


Client



Amazon Elastic Cloud Compute (Amazon EC2)

This is the core of the Amazon Cloud,



on
ices™
usage

AWS Management Console - Mozilla Firefox

File Edit View History Bookmarks Tools Help

amazon.com https://console.aws.amazon.com/ec2/home#?c=EC2&s=Instances

AWS Management Console

aws.amazon.com AWS | Products | Developers | Community | Support | Account

Welcome, Dr William Buchanan | Settings | Sign Out

Amazon EC2 Amazon Elastic MapReduce Amazon CloudFront

Navigation

Region: US East

EC2 Dashboard

INSTANCES

- Instances
- Spot Requests

IMAGES

- AMIs
- Bundle Tasks

ELASTIC BLOCK STORE

- Volumes
- Snapshots

NETWORKING & SECURITY

- Elastic IPs
- Security Groups
- Key Pairs
- Load Balancers

My Instances

Launch Instance Instance Actions Reserved Instances Show/Hide Refresh Help

Viewing: All Instances All Instance Types 1 to 1 of 1 Instances

Instance	AMI ID	Root Device Type	Type	Status	Lifecycle	Public DNS	Security Groups	Key Pair Name
<input checked="" type="checkbox"/> i-0d895566	ami-45c22e2c	ebs	m1.small	running	normal	ec2-204-236-199-96.compute-1.amazonaws.com	bill	billwindows2

EC2 Instance: i-0d895566

Description Monitoring

AMI ID:	ami-45c22e2c	Zone:	us-east-1c
Security Groups:	bill	Type:	m1.small
Status:	running	Owner:	103269750866
Reservation:	r-fcd18794	Ramdisk ID:	-
Platform:	windows	Key Pair Name:	billwindows2008
Kernel ID:	-	Monitoring:	enabled
AMI Launch Index:	0	Elastic IP:	-
Root Device:	/dev/sda1	Root Device Type:	ebs
Block Devices:	/dev/sda1=vol-49e43920:attached:2010-03-07T19:58:20.000Z:true		
Lifecycle:	normal		
Public DNS:	ec2-204-236-199-96.compute-1.amazonaws.com		

© 2008 - 2009, Amazon Web Services LLC or its affiliates. All right reserved. Feedback Support Privacy Policy Terms of Use An amazon.com company



Client

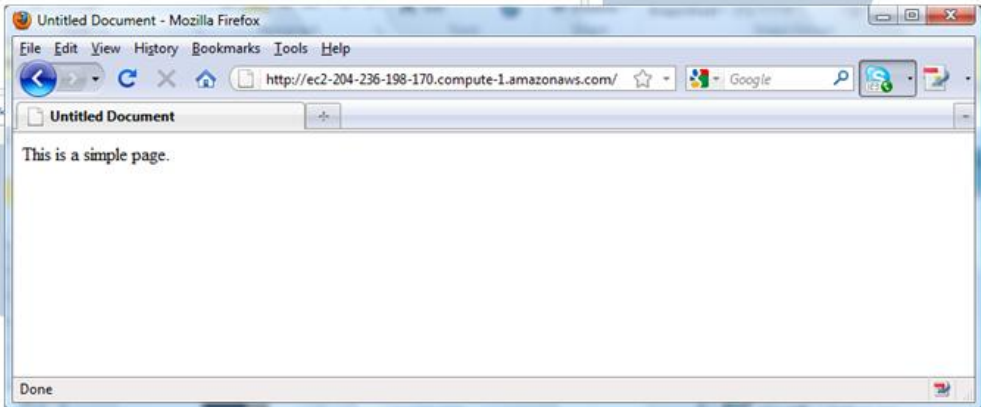


Amazon Elastic Cloud Compute (Amazon EC2)



Amazon
Services™
usage

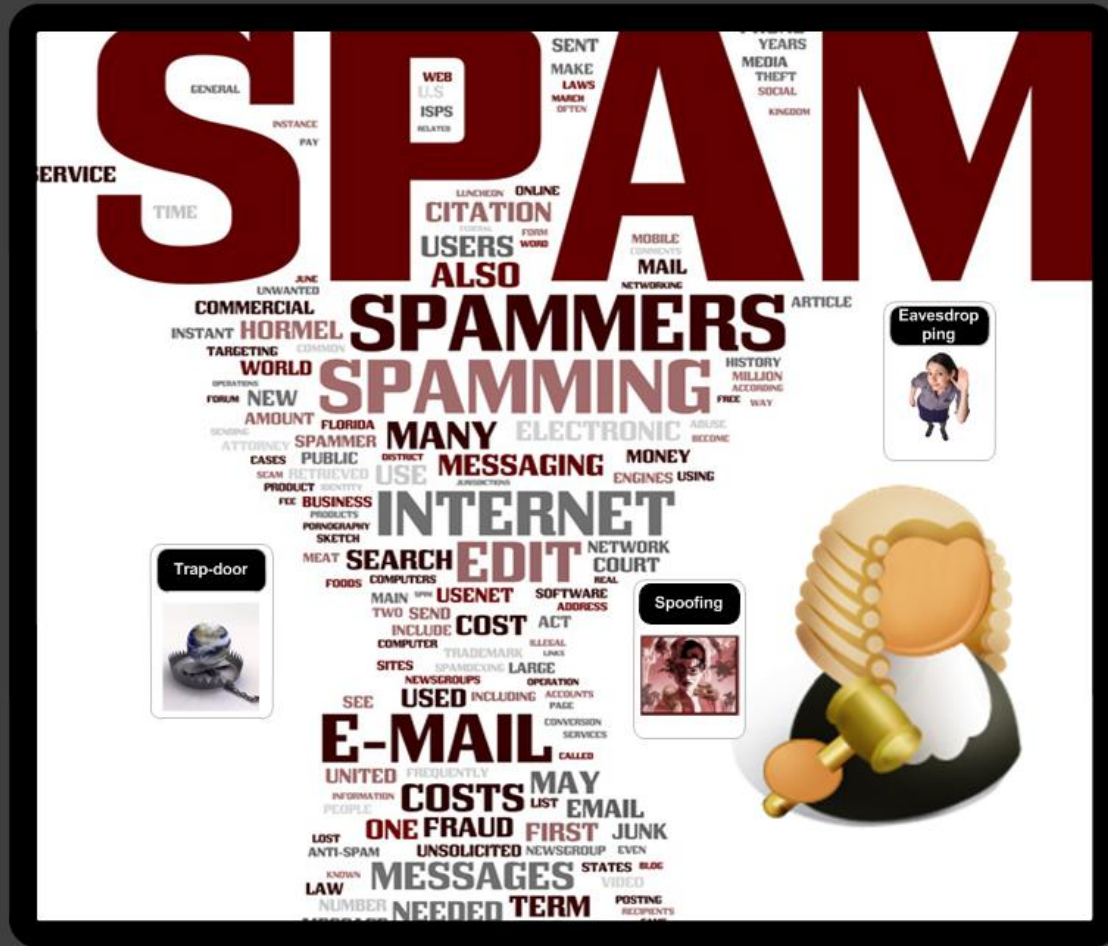
Connection Method	Protocol	From Port	To Port	Source (IP or group)	Actions
RDP	tcp	3389	3389	0.0.0.0/0	Remove
HTTP	TCP	80	80	0.0.0.0/0	Save



AWS

Cloud

Cloud Computing



Amazon EC2 Programming

Amazon Elastic Compute Cloud
Command Line Tools Reference (API Version 2009-11-30)

Welcome
What's New
Getting the Command Line Tools
AMI Tools Reference
API Command Line Tools Reference

ec2-describe-images

Description

Returns information about AMIs, AKIs, and ARIs. This includes image type, product codes, architecture, and kernel and RAM disk IDs. Images available to you include public images, private images that you own, and private images owned by other users for which you have explicit launch permissions.

Launch permissions fall into three categories:

Launch Permission	Description
public	The owner of the AMI granted launch permissions for the AMI to the <code>all</code> group. All users have launch permissions for these AMIs.
explicit	The owner of the AMI granted launch permissions to a specific user.
implicit	A user has implicit launch permissions for all AMIs he or she owns.

The list of AMIs returned can be modified by specifying AMI IDs, AMI owners, or users with launch permissions. If no options are specified, Amazon EC2 returns all AMIs for which the user has launch permissions.

If you specify one or more AMI IDs, only AMIs that have the specified IDs are returned. If you specify an invalid AMI ID, a fault is returned. If you specify an AMI ID for which you do not have access, it will not be included in the returned results.

If you specify one or more AMI owners, only AMIs from the specified owners and for which you have access are returned. The results can include the account IDs of the specified owners, `amazon` for AMIs owned by Amazon or `self` for AMIs that you own.


If you specify a list of executable users, only users that have launch permissions for the AMIs are returned. You can specify account IDs (if you own the AMI(s)), `self` for AMIs for which you own or have explicit permissions, or `all` for public AMIs.



i-0d895567

Machine Instances

- Command line tools
- Web Service Console



i-0d895566

Amazon Machine Images (AMI)

Return images owned by a specific owner

```
$ ec2-describe-images -o AKIAIWUMTTAZYST2I2AA
IMAGE ami-45c22e2c powerdns/image.manifest.xml 495219933132 available private
$ ec2-run-instances ami-45c22e2c -n 10
```

AWS

Cloud

Cloud/Grid Computing

- Provide an introduction to cluster, grid and cloud infrastructures.
- Define an example of grid computing, and its advantages.
- Show an example of using a Cloud Infrastructure.
- Define the usage of Amazon Web Services, including S3 and the EC2 Cloud.
- Outline the usage of RESTful Web Services.

