CC Lecture 8b - Security

Open Data Management & the Cloud

(Data Science & Scientific Computing / UniTS – DMG)

Outline

□ Introduction

□ Infrastructure security

□ Data security

□ Identity and access management

Introduction

- What should I do when I use the cloud and I want some privacy?
- Many security problems in non-cloud environment are still applicable

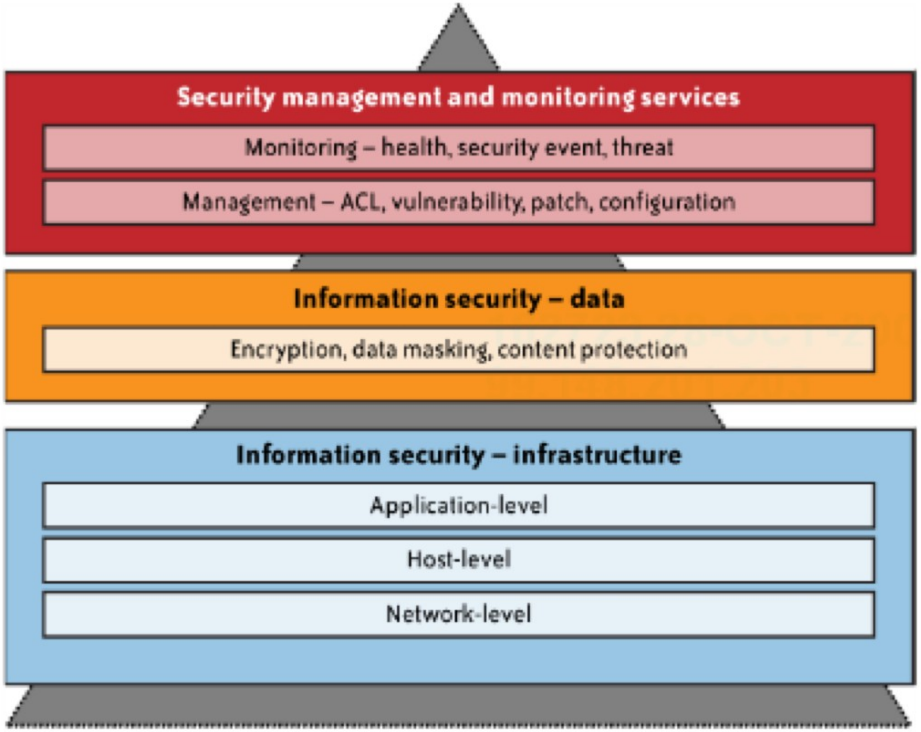

Security management and monitoring services

Monitoring – health, security event, threat

Management – ACL, vulnerability, patch, configuration

Information security – data

Encryption, data masking, content protection

Information security – infrastructure

Application-level

Host-level

Network-level

Main user concerns

- Availability of cloud services
- Third-party control
- Insecure APIs
- Data loss or leakage
- Account or service hijacking

- Unauthorized Access: Trust.

General considerations

Security boundary separating the client’s and vendor’s responsibilities.

- Different types of cloud computing service models provide different levels of security
- Focus on public clouds
- Different levels:
  - Network level
  - Host level
  - Application level

The XaaS boundaries

Service Layers Definition

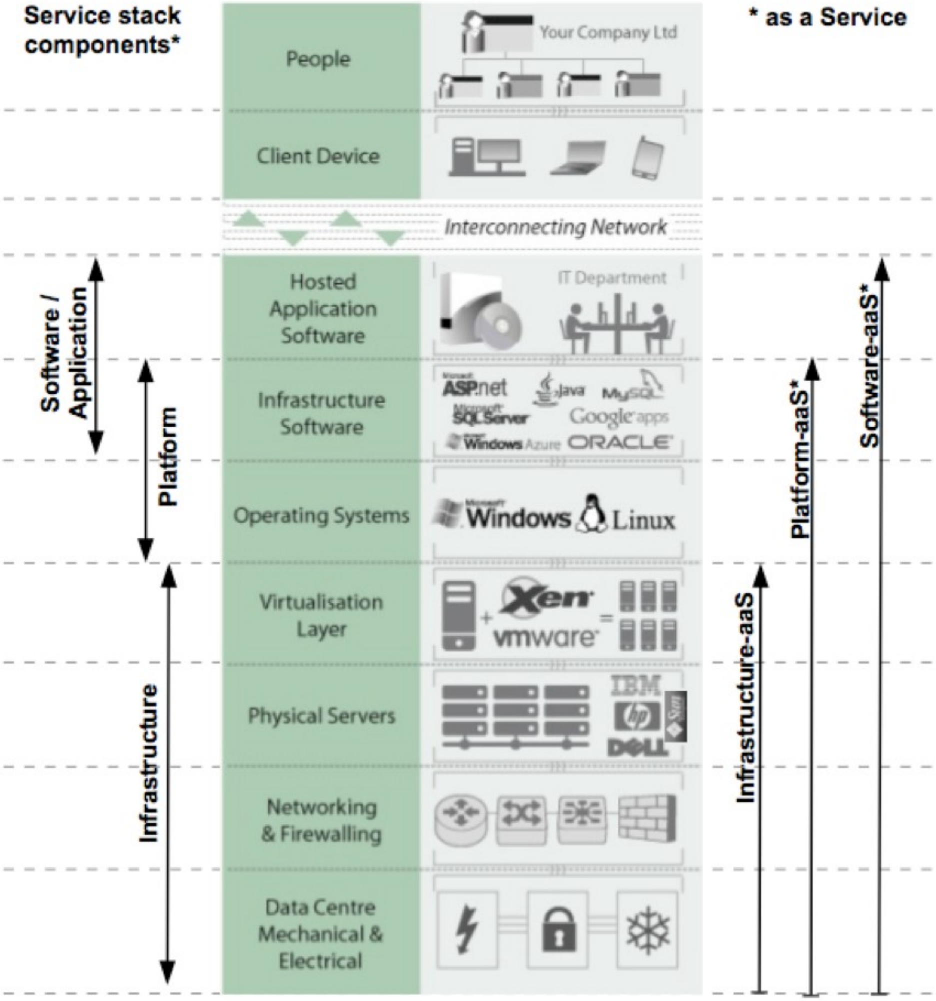
Service stack components*

* as a Service

People

Your Company Ltd

Client Device

Interconnecting Network

Hosted Application Software

IT Department

Infrastructure Software

ASP.net

Java

MySQL

SQL Server

Google apps

Windows Azure

ORACLE

Operating Systems

Windows

Linux

Virtualisation Layer

Xen

vmware

Physical Servers

IBM

hp

Sun

DELL

Networking & Firewalling

Data Centre Mechanical & Electrical

Software / Application

Platform

Infrastructure

Infrastructure-aaS

Platform-aaS*

Software-aaS*

Notes:
Brand names for illustrative / example purposes only, and examples are not exhaustive.

** Assumed to incorporate subordinate layers.*

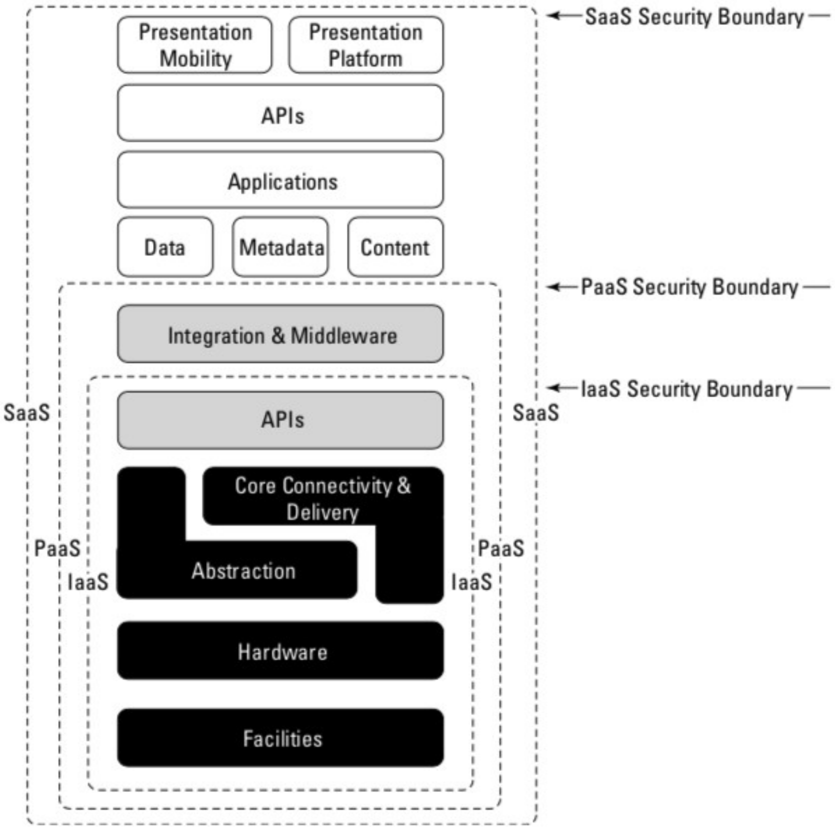
SaaS Security Boundary

Presentation Mobility

Presentation Platform

APIs

Applications

Data

Metadata

Content

PaaS Security Boundary

Integration & Middleware

IaaS Security Boundary

SaaS

SaaS

APIs

Core Connectivity & Delivery

PaaS

PaaS

IaaS

IaaS

Abstraction

Hardware

Facilities

Security Boundary

- Who has responsibility for a particular security mechanism?
- Where is the boundary between the responsibility of the service provider and the responsibility of the customer?
- Cloud Security Alliance: CSA is an industry working group that studies security issues in cloud computing and offers recommendations to its members.

The top concern for cloud users

- Security is the top concern for cloud users;
- Top concerns for users are:
  - Unauthorized access to confidential information and data theft;
  - User control over the life cycle of data;
  - Lack of standardization;
  - Uncontrolled technology evolution;
  - Multitenancy
  - Legal framework;

Privacy and Trust

The term privacy refers to the right of an individual, a group of individuals, or an organization to keep information of a personal or proprietary nature from being disclosed to others.

Trust in the context of cloud computing is intimately related to the general problem of trust in online activities.

“trust of a party A to a party B for a service X is the measurable belief of A in that B behaves dependably for a specified period within a specified context (in relation to service X)”

The platform security level is reduced to the security level of the most vulnerable application running on the platform. (isolate with VM and containers)

Network Level Security

- confidentiality and integrity of data-in-transit
  - Amazon had security bugs with digital signature on SimpleDB, EC2, and SQS accesses (in 2008)
- Less or no system logging /monitoring
  - Only cloud provider has this capability
  - Thus, difficult to trace attacks
- Reassigned IP address
  - Expose services unexpectedly
  - spammers using EC2 are difficult to identify
- Availability of cloud resources
  - Some factors, such as DNS, controlled by the cloud provider.
- Physically separated tiers become logically separated
  - E.g., 3 tier web applications

OS Security

- Implement security on: Access control, authentication usage, and cryptographic usage policies are all elements of mandatory OS security.

- Applications with special privileges that perform security-related functions are called trusted applications. Such applications should only be allowed the lowest level of privileges required to perform their functions.

- An OS poorly isolates one application from another, and once an application is compromised, the entire physical platform and all applications running on it can be affected.

Host Level

- Hypervisor security
  - “zero-day vulnerability” in VM, if the attacker controls hypervisor
- Virtual machine security
  - ssh private keys (if mode is not appropriately set)
  - VM images (especially private VMs)
  - Vulnerable Services

XaaS and VM security advantages

- Simple way to implement resource management policies
- Improved intrusion prevention and detection.
- Secure logging and intrusion protection.
- More efficient and flexible software testing

But….

VM and XaaS Risks

- Explosion of number of VMs
- Snapshots (roll back to state that can be exploited)
- Shared Images
  - backdoors and leftover credentials,
  - unsolicited connections,
  - Malware
- 22% of the scanned Linux AMIs contained credentials allowing an intruder to remotely log into the system.

Identity and Access Management

- IAM components
  - Authentication
  - Authorization
  - Auditing
- IAM processes
  - User management
  - Authentication management
  - Authorization management
  - Access management & access control
  - Propagation of identity to resources
  - Monitoring and auditing

Federated Authentication

The main purpose of federated identity management is to allow registered users of a certain domain to access information from other domains in a smooth way without having to provide any extra administrative user information

- Gives a delegated mechanism to manage user identification among different entities and within different subjects
- Provides a set of attributes to an authenticated users to be used by the final application.

Federated Authentication features

Identity Provider asserts authentication and identity information about users. Keep your credential at your institute/company.

Service Providers check and consume this information for authorization and make it available to an application.

Reduce Work

Protects User Information

Provides current A&A info

Insulate from service compromise

Role of federations

A group of organizations running IdPs and SPs that agree on a

common set of rules and standards

Based on TRUST!

Defines agreements and rules

Operates discovery services

An organization may belong to more than one federation

Storage security example

Any data can be intercepted and modified in a WAN.

- Implement data encryption;
- Implement secure authentication and Authorization
- Implement a brokered data access;

Technologies

OAuth (Open Authentication)

Security Assertion Markup Language (Shibboleth)

OpenID

Unity (solution for identity, federation and inter-federation management)

SAML example

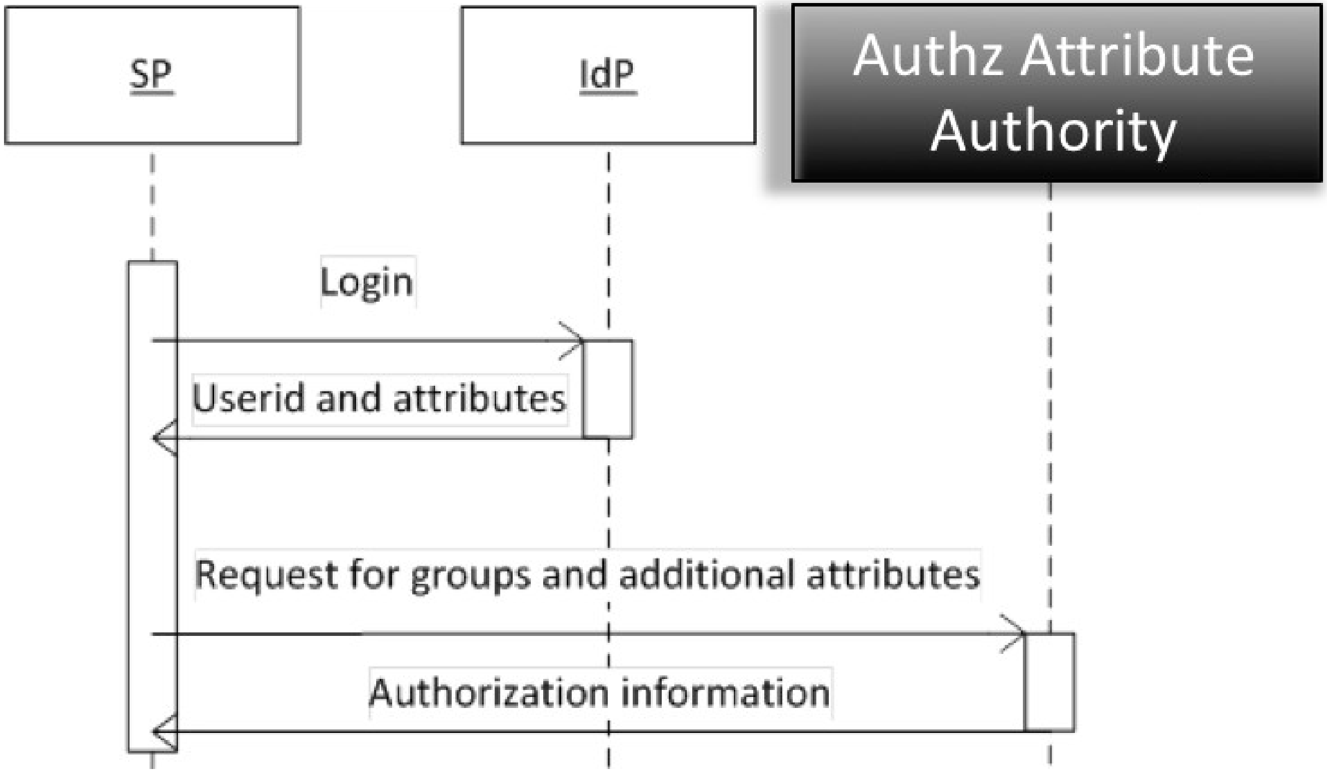
SP
IdP
Authz Attribute Authority
Login
Userid and attributes
Request for groups and additional attributes
Authorization information

Attributed based

- TERENA security policy and security model
- eduPerson schema for LDAP and....SAML
- Attribute based authorization

```
urn:schac:userStatus:au:uq.edu.au:service:mail:receive:disabled
```

Attribute
Value

EduGain

The eduGAIN service interconnects identity federations around the World, simplifying access to content, services and resources for the global research and education community. eduGAIN enables the trustworthy exchange of information related to identity, authentication and authorisation (AAI).

Applied Security: Storage

- The single largest security concern that most organizations should have.
- As with any WAN traffic, any data can be intercepted and modified.
- Data can be located anywhere in the cloud provider data centers
- Data can be accessed by provider personnel.

How to protect data

- Identify the security boundary separating the client’s and vendor’s responsibilitie
- Determine the sensitivity of the data to risk
- Data should be transferred and stored in an encrypted format.
- Separate clients from direct access to shared cloud storage.

Data segregation

- Isolate data from direct client access creating a layered access to the data.
- Data segregation based on tenants

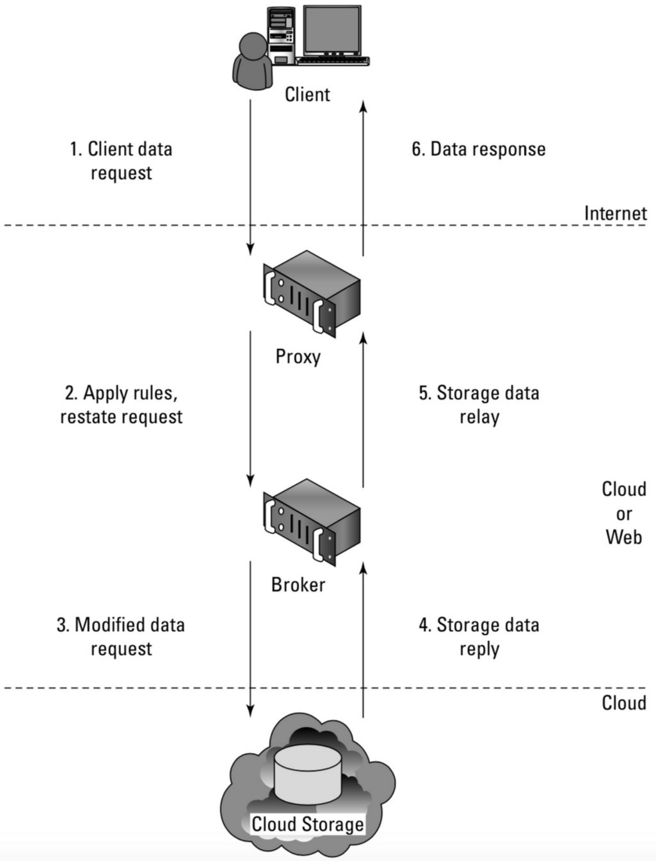

Encryption

- Most cloud service providers store data in an encrypted form (e.g. Amazon S3 256-bit Advanced Encryption Standard) on server side or client side.
- Some example of java code here:

  https://docs.aws.amazon.com/AmazonS3/latest/dev/UsingClientSideEncryption.html
- Problems:
  - a problem with encrypted data may result with data that may not be recoverable.
  - it does nothing to prevent data loss: keep you keys!!!!!

Before moving into the Cloud

- Determine which resources (data, services, or applications) you are planning to move to the cloud.
- Determine the sensitivity of the resource to risk.
    - Risks that need to be evaluated are loss of privacy, unauthorized access by others, loss of data, and interruptions in availability.
- Determine the risk associated with the particular cloud type for a resource
- Take into account the particular cloud service model that you will be using.
- If you have selected a particular cloud service provider, you need to evaluate its system to understand how data is transferred, where it is stored, and how to move data both in and out of the cloud.