Security Engineering

Lecture 18 – Clouds and all that

Fabio Massacci

NIST Definition

- ***A model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.***

Cloud Architectural Solutions

- ***SaaS (Software as a Service)***
  - A provider licenses an application to customers for use as a servi on demand.
  - vendors host application on own web servers or download the application to consumer device, disabling it after contract expires.
- ***PaaS (Platform as a service)***
  - delivery of computing platform & solution stack as a service.
  - facilitates deployment of applications without cost & complexity of buying and managing hardware & software layers.
  - Environment supports lifecycle for building & running applications
- ***IaaS (Infrastructure as a Service)***
  - delivery of computer infrastructure as a service typically a virtualized environment managed in an integrated and efficient way.
  - Offers computing as a service billed on a utility basis and amount of resources consumed

World Scale Computing aka Cloud

- ***Extremely large datacenters 1,000's to 10,000's of commodity computers***
  - 5-7x cheaper than provisioning a medium-sized (100's machines) facility of high performance computers
  - Build-out driven by demand growth (more users)
  - More reliable: hardware failures, DDoS, etc
- ***Enabling factors***
  - Available connection by broadband internet
  - x86 as universal ISA, fast virtualization
  - Infrastructure software: eg Hadoop, Google FileSystem
  - Standard software stack
- ***Yet this is an old idea (more than 20yrs old)***
  - ASP = Application Service Providers
  - Software hosted in the infrastructure vs. installed on local servers or desktops
  - Why now and not then?

Reasons for Cloud Virtualization

- ***Server consolidation (Physical-to-Virtual (P2V) transformation)***
  - many small physical servers → one larger physical server, to increase utilization of hw
  - The large server can "host" many such "guest" virtual machines
- ***Inspection and isolation***
  - A virtual machine can be more easily controlled and inspected from outside than a physical one, and its configuration is more flexible.
- ***Provisioning and relocation***
  - A new VM can be provisioned as needed without the need for an up-front hardware purchase.
  - VM can easily be relocated from one physical machine to another as needed.
- ***Disaster recovery scenarios***
  - Because of easy relocation
  - ONLY work with machines in different locations. If you only have one big server won't work
- ***But which is the real reason?***

Capacity Planning

- Capacity must exceed maximum demand if you want to meet demand at all times

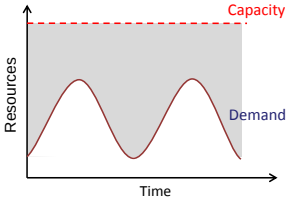

Unused resources

Static data center

Risks of underprovisioning

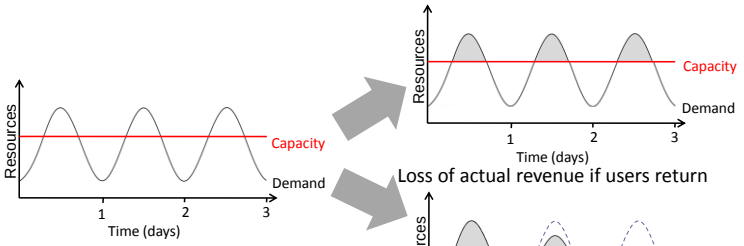

Loss of actual revenue if users return

Loss of actual revenue AND future users

Static vs Dynamic Provisioning

- Static provisioning for peak: wasteful, but necessary for SLA

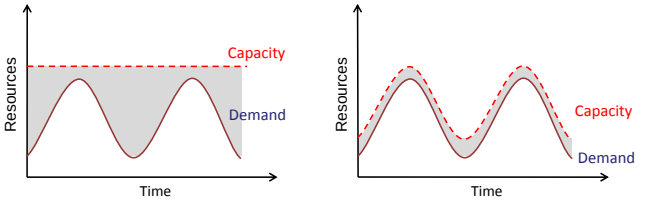

"Statically provisioned" data center

"Virtual" data center in the cloud

Unused resources

ASP vs Cloud Computing

- ***The key idea of ASP – 1989***
  - Eastman Kodak decided to outsource the IT system that underpinned their entire enterprise to IBM
- ***The key idea of Cloud Computing***
  - Animoto traffic doubled every 12 hours for 3 days when released as Facebook plug-in
  - Scaled from 50 to >3500 servers
  - ...then scaled back down
- ***What is the difference?***

ASP vs Cloud Computing

- ***The key idea of ASP***
  - In 1989 Eastman Kodak outsource the whole IT system to IBM
- ***The key idea of Cloud Computing***
  - Animoto traffic doubled every 12 hours for 3 days when released as Facebook plug-in. From 50 to >3500 servers then scaled back down
  - Thales French Tax Portal, half a million income tax returns filed on-line in the last few days before the close of the tax year, then nothing happens for a year
- ***What is the difference?***
  - Kodak IT users were known in advance and they didn't change.
  - Kodak Co. didn't want to pay upfront for the hardware cost or the IT expertise
  - Animoto/Thales users not known in advance but their "work" known and standard.
  - Animoto/Thales users came from the internet! Without "internet users", cloud computing doesn't make sense
    - indeed networking and data transfer costs are where cloud providers rip you off

Classifying Clouds

- ***IaaS: Amazon EC2, Rackspace***
- ***PaaS: Microsoft Azure***
- ***SaaS: GoogleApps, Force.com, SAP ByD***
- ***Tradeoff:***
  - flexibility/portability vs. "built in" functionality

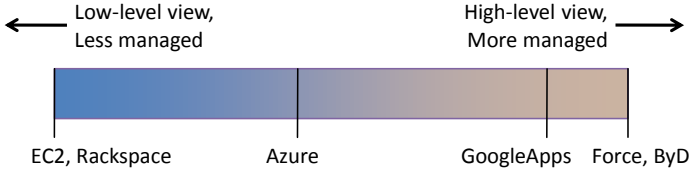

Sample "Cloud" Scenarios

- ***Running one or more applications not supported by host OS***
  - A virtual machine running required guest OS could allow the desired applications to be run
- ***Evaluating an alternate operating system***
  - The new OS could be run within a VM
- ***Server virtualization***
  - Multiple virtual servers could be run on a single physical server, more fully utilize the hardware of the physical server.
- ***Duplicating specific environments***
  - A VM could be duplicated and installed on multiple hosts.
- ***Creating a protected environment***
  - If guest OS running on a VM becomes infected with malware, host operating system's exposure may be limited)

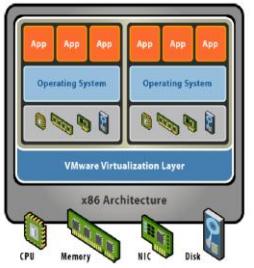

vmware Parallels Xen

Source: Wikipedia, VMWare

EC2 configuration

- ***EC2 uses Xen, with up to 8 instances per physical machine***

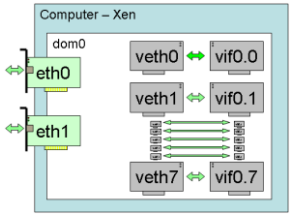

Dom0 is the first instance on the machine, connected to physical adapter

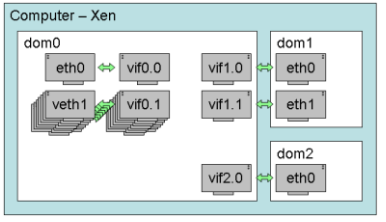

All other instances route to external world via dom0

[Figures from Xen Wiki by R. Hasan]

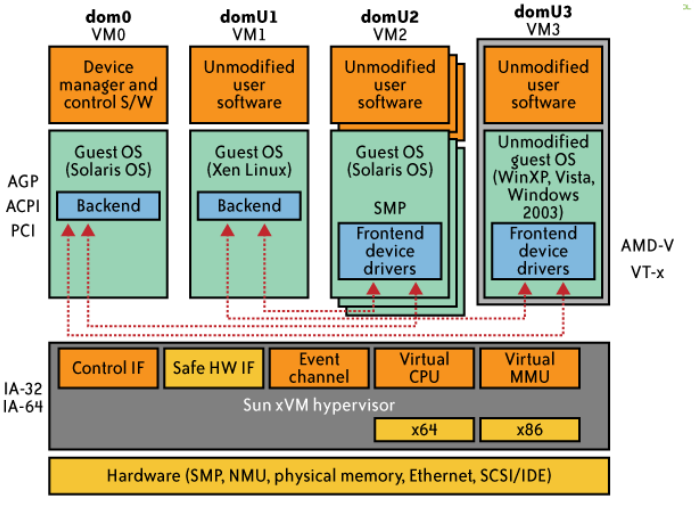

From [6] Cloud Security and Privacy by Mather and Kumaraswamy

Top Customers Concerns (IBM Survey 2010)

- Protection of intellectual property and data — 30%
- Ability to enforce regulatory or contractual obligations — 21%
- Unauthorized use of data — 15%
- Confidentiality of data — 12%
- Availability of data — 9%
- Integrity of data — 8%
- Ability to test or audit a provider's environment — 6%
- Other — 3%

External Attackers

- ***Potential Objectives***
  - Network analysis/tampering
  - Application Data exfiltration/tampering (eg web defacing)
  - Control of VM
  - Escape from VM…
  - "Cartography"
- ***What can they do?***
  - Listen to or insert into network traffic
  - Compromise WebApp
  - Install malware into VM
  - Launch (D)DoS
  - Probe cloud structure
  - Scrap data from VM and HV

Malicious Insiders

- ***At client***
  - Learn passwords/authentication information
  - Gain control of VMs
- ***At cloud provider***
  - Log client communication, monitor network communication, application patterns
  - Read unencrypted data
  - Peek into VMs, or make copies of VMs
- ***Possible Objectives***
  - Gain information about client data/behavior
  - Sell information or use for himself
    - Google's Engineer David Barksdale fired in 2010
- ***Difference with standard attacks?***

More attackers

- ***Sloppy insiders***
  - Regulatory compliance is enforced but with local (lower)standards or no standards
    - E.g. Ireland vs Germany for privacy legislation
  - The legal responsibility stays with the customer
    - If the cloud provider is sloppy the customers pays the legal bills
- ***Curious law enforcement officers***
  - Obviously not "your" law enforcement guys
  - Guys from country where cloud providers are located can ask to access data
    - Exists also the other coin of the medal: LEOs from customers' country wanting access to the data….
- ***Heat, cold, earthquakes, and other man-made failures***

Malware and VM

- ***Normal attacks → Compromise the VM***
- ***VM specific attacks***
  - Simplest → detect it.
    - As security researchers rely on VM emulators, malicious code alter their behavior (including refusing to run) if a VM emulator is detected
  - Clever → take-over and look for memory leaks of "siblings"
    - Guest OS doesn't know is running on a VME so it does not "clean" memory after use
    - Hypervisor doesn't do because it assumes VM wants to do whatever with the hw
    - Several VM co-exhist on the same physical machines…
  - Harsher → denial-of-service
    - this type of attack causes the virtual machine emulator to exit.
  - Most interesting → escape from its protected environment
    - Take over the actual hypervisor
    - VENOM, CVE-2015-3456, is a vulnerability in the open source virtual floppy drive code used by many VME. An attacker can escape from the confines of an affected virtual machine (VM) guest and potentially obtain code-execution access to the host.

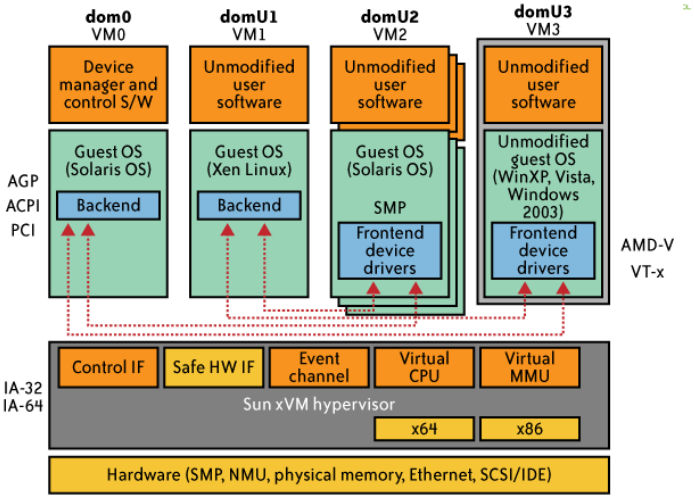

From [6] Cloud Security and Privacy by Mather and Kumaraswamy

The Insider Attack

- ***The biggest problem for the insider attack is how to select a "good" provider***
  - Legislation
  - Type of security measures
  - Transparency of security measures (certification etc.)
- ***Cloud Security Alliance***
  - List different providers and different measures they implemenent
    - China Telecom answered yes to all questions, including "is you data center located in place with potential natural disasters"
- ***Specific exercise on this one***

From ASP to Multi-Tenancy

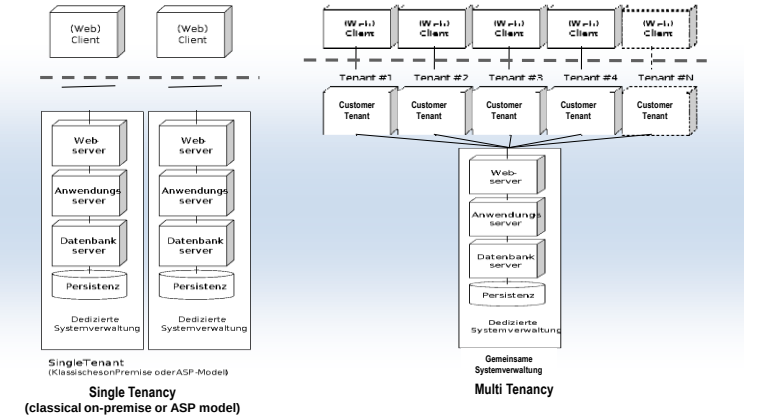

Source: SAP

Integration of WS+DB+OS I

- ***Observation***
  - a DBMS has to stop tenants from interfering with each other, and with the DBMS
  - DBMS same tasks of O.S.
- ***To avoid duplication, could give these tasks to O.S.***
  - DBMS runs as a set of O.S. processes;
  - System processes for general database management tasks
  - Each DB tenant is mapped to a separate O.S. process.
- ***Operating system can distinguish between users***
  - if each database object is stored in its own file, O.S. can do all the access control
- ***DBMS only translates user queries → the "Apache Model"***

Integrating WS+DB+OS II

- ***Allocating an individual operating system process to every database tenant wastes memory resources and does not scale up to many users of many tenants***
- ***Letting processes handle the database requests of several tenants and their users***
  - saves memory
  - but the DBMS becomes responsible for access control
- ***Similar considerations for storing database objects***
  - for small objects, having a separate file for each object is wasteful
  - If the operating system does not control access of database tenants, several database objects can be collected in one operating system file

Efficient & Sclabale Multi-Tenancy

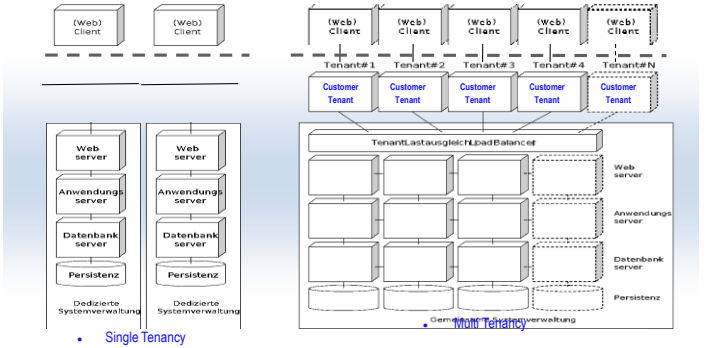

- Single Tenancy
- (classical on-premise or ASP model)

- Multi Tenancy

Source: SAP

The less isolation the "better"…

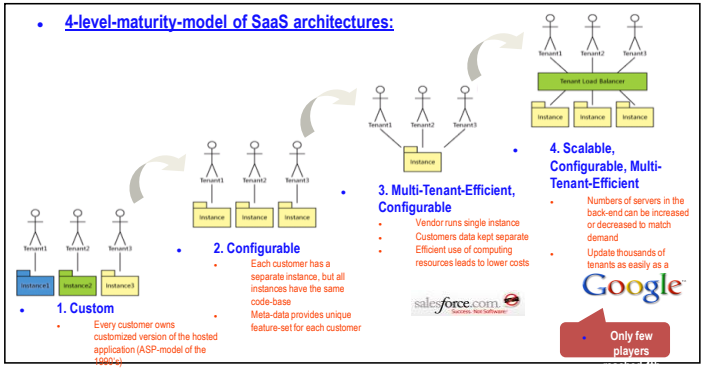

[MSDN, F. Chong and G. Carraro, "Architecture Strategies for Catching the Long Tail", http://msdn.microsoft.com/en-us/library/aa479069.aspx, April 2006.]

Performance wins again

- ***The hit of crossing the kernel/OS boundary:***
  - Original Apache implementation forked a process to run each CGI:
  - Could attenuate file access for sub-process
  - protected memory/data of server from rogue script
    - Very close to least privilege
- ***Too expensive for***
  - a small script (fork, exec, copy data to/from the server process), etc.
  - if this is repeated millions or billions of times…
    - can have more hardware but hardware don't scale equally well than clients
    - and you started all that to avoid having as much hardware as clients…
- ***current push is to run the scripts in the server.***
  - See Node.JS raison d'etre…
  - Throw out least privilege
- ***Similar situation with DBs, web browsers, file systems, etc.***

Additional readings

- ***Gollmann – Computer Security***
  - Ch. 8 – Operating Systems
  - Ch. 9 – Databases
- ***NIST Guide on Hypervisor***
  - csrc.nist.gov/publications/nistpubs/800-125/SP800-125-final.pdf
- ***Barroso & Hölzle, "The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines", 2009***