cloud security

EECE 571B “Computer Security”

Konstantin Beznosov

overview of cloud computing


software
hardware
network
facilities
IT labor
management tools
power/cooling
support
maintenance
security
disaster recovery
backup
Acquisition cost is 10% of IT Spend
Operating cost is 90% of IT Spend
Source: IDC

basics

- definition
  - a collection/group of integrated and networked hardware, software and Internet infrastructure (called a platform)
- properties/characteristics
  - remotely hosted: services or data are hosted on remote infrastructure
  - abstracted: cloud platforms hide the complexity and details of the underlying infrastructure from users and applications by providing very simple graphical interface or API
  - ubiquitous: on demand services, that are always on, anywhere, anytime and any place
  - commodified: utility computing model similar to that of traditional utilities, like gas and electricity
    - pay for use and as needed, elastic (scale up and down in capacity and functionalities).
    - available to the general public, organizations, and companies

Cloud Architecture

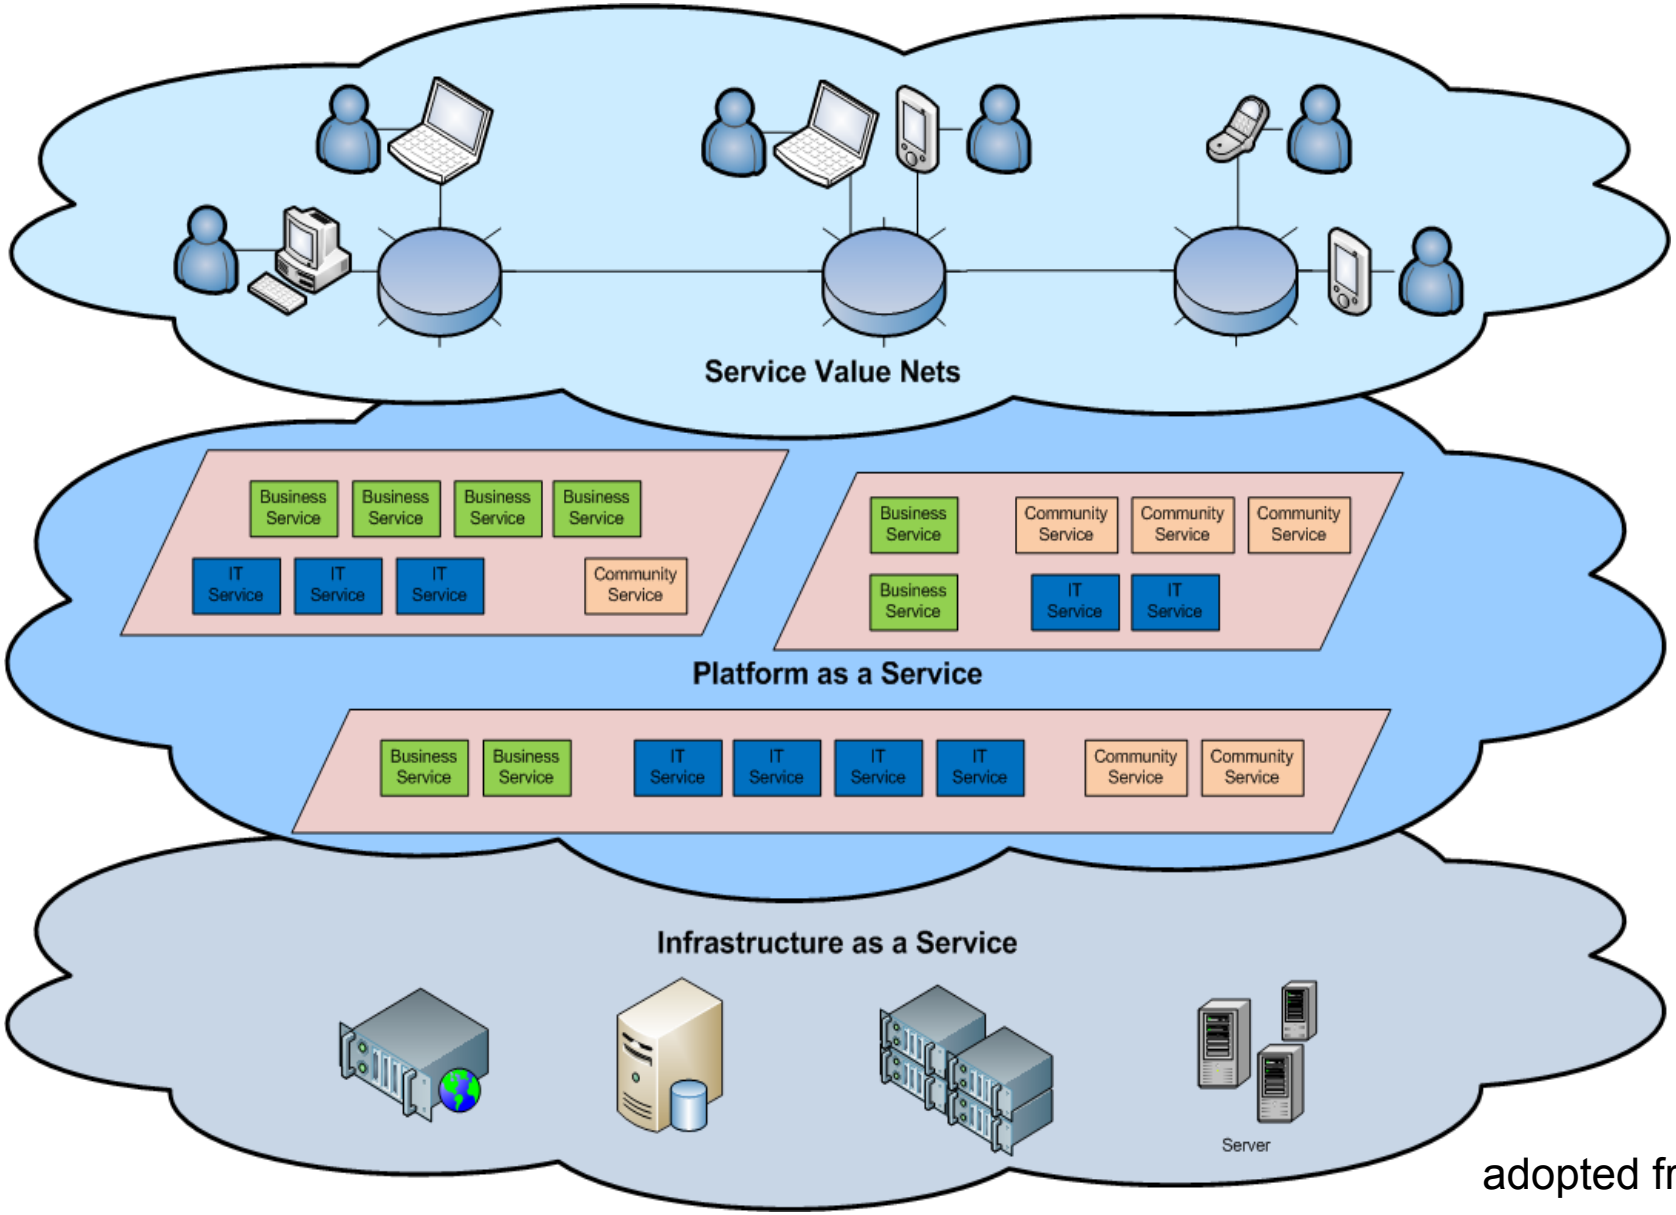

adopted from [1]

Different Cloud Computing Layers

Application Service (SaaS)	MS Live/ExchangeLabs, IBM, Google Apps; Salesforce.com Quicken Online, Zoho, Cisco
Application Platform	Google App Engine, Mosso, Force.com, Engine Yard, Facebook, Heroku, AWS
Server Platform	3Tera, EC2, SliceHost, GoGrid, RightScale, Linode
Storage Platform	Amazon S3, Dell, Apple, ...

adopted from [1]

Cloud Computing Service Layers

	Services	Description
Application Focused	Services	Services – Complete business services such as PayPal, OpenID, OAuth, Google Maps, Alexa
	Application	Application – Cloud based software that eliminates the need for local installation such as Google Apps, Microsoft Online
	Development	Development – Software development platforms used to build custom cloud based applications (PAAS & SAAS) such as SalesForce
Infrastructure Focused	Platform	Platform – Cloud based platforms, typically provided using virtualization, such as Amazon ECC, Sun Grid
	Storage	Storage – Data storage or cloud based NAS such as CTERA, iDisk, CloudNAS
	Hosting	Hosting – Physical data centers such as those run by IBM, HP, NaviSite, etc.

adopted from [1]

CLOUD TAXONOMY

Infrastructure Services

Storage
- Amazon S3
- Amazon EBS
- CTERA Portal
- Mosso Cloud Files
- Nirvanix

Compute
- Amazon EC2
- Serve Path GoGrid
- Elastra
- Mosso Cloud Servers
- Joyent Accelerators
- AppNexus
- Flexiscale
- Elastichosts
- Hosting.com CloudNine
- Terramark
- GridLayer
- iTRiCiTY
- LayeredTech

Services Management
- RightScale
- enStratus
- Scalr
- CohesiveFT
- Kaavo
- CloudStatus
- Ylastic
- Dynect
- CloudFoundry
- NewRelic
- Cloud42

Cloud Software

Data
- 10Gen MongoDB
- Oracle Coherence
- Gemstone Gemfire
- Apache CouchDb
- Apache HBase
- Hypertable
- TerraCotta
- Tokyo Cabinet
- Cassandra
- memcached

Compute
- Globus Toolkit
- Xeround
- Beowulf
- Sun Grid Engine
- Hadoop
- OpenCloud
- Gigaspaces
- DataSynapse
- Xeround

Cloud Management
- 3Tera App Logic
- OpenNebula
- Open.ControlTier
- Enomaly Enomalism
- Altor Networks
- VMware vSphere
- OnPathTech
- CohesiveFT VPN Cubed
- Hyperic
- Eucalyptus
- Reductive Lbs Puppet
- OpenQRM
- Appistry

Appliances
- PingIdentity
- Symplified
- rPath
- Vordel

File Storage
- EMC Atmos
- ParaScale
- Zmamda
- CTERA

Platform Services

General Purpose
- Force.com
- Etelos
- LongJump
- AppJet
- Rollbase
- Bungee Labs Connect
- Google App Engine
- Engine Yard
- Caspio
- Qrimp
- MS Azure Services Platform
- Mosso Cloud Sites

Business Intelligence
- Aster DB
- Quantivo
- Cloud9 Analytics
- Blink Logic
- K2 Analytics
- LogiXML
- Oco
- Panorama
- PivotLink
- Sterna
- ColdLight Neuron
- Infobright
- Vertica

Integration
- Amazon SQS
- MuleSource Mule OnDemand
- Boomi
- SnapLogic
- OpSource Connect
- Cast Iron
- Microsoft BizTalk Services
- gnip
- SnapLogic SaaS Solution Packs
- Appian Anywhere
- HubSpan
- Informatica On-Demand

Development & Testing
- Keynote Systems
- Mercury
- SOASTA
- SkyTap
- Aptana
- LoadStorm
- Collabnet
- Dynamsoft

Database
- Google BigTable
- Amazon SimpleDB
- FathomDB
- Microsoft SDS

OpenCrowd

Software Services

Billing
- Aria Systems
- eVapt
- OpSource
- Redi2
- Zuora

Financials
- Concur
- Xero
- Workday
- Beam4d

Legal
- DirectLaw
- Advologix
- Fios
- Sertifi

Sales
- Xactly
- LucidEra
- StreetSmarts
- Success Metrics

Desktop Productivity
- Zoho
- IBM Lotus Live
- Google Apps
- Desktoptwo
- Parallels
- ClusterSeven

Human Resources
- Taleo
- Workday
- iCIMS

Content Management
- Clickability
- SpringCM
- CrownPoint

Backup & Recovery
- JungleDisk
- Mozy
- Zmanda Cloud Backup
- OpenRSM
- Syncplicity

CRM
- NetSuite
- Parature
- Responsys
- Rightnow
- Salesforce.com
- LiveOps
- MSDynamics
- Oracle On Demand

Document Management
- NetDocuments
- Questys
- DocLanding
- Aconex
- Xythos
- Knowledge TreeLive
- SpringCM

Collaboration
- Box.net
- DropBox

Social Networks
- Ning
- Zembly
- Amitive

Updated as of May 4, 2009

Virtualization

- Virtual workspaces:
  - An abstraction of an execution environment that can be made dynamically available to authorised clients by using well-defined protocols,
  - Resource quota (e.g. CPU, memory share),
  - Software configuration (e.g. O/S, provided services).
- Implement on Virtual Machines (VMs):
  - Abstraction of a physical host machine,
  - Hypervisor intercepts and emulates instructions from VMs, and allows management of VMs,
  - VMWare, Xen, etc.
- Provide infrastructure API:
  - Plug-ins to hardware/support structures

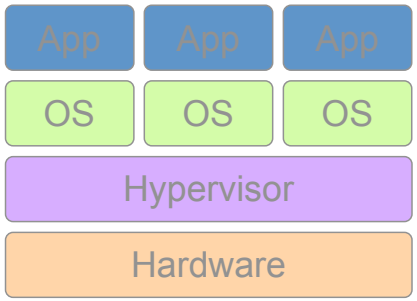

Virtualized Stack

Virtual Machines

- VM technology allows multiple virtual machines to run on a single physical machine.

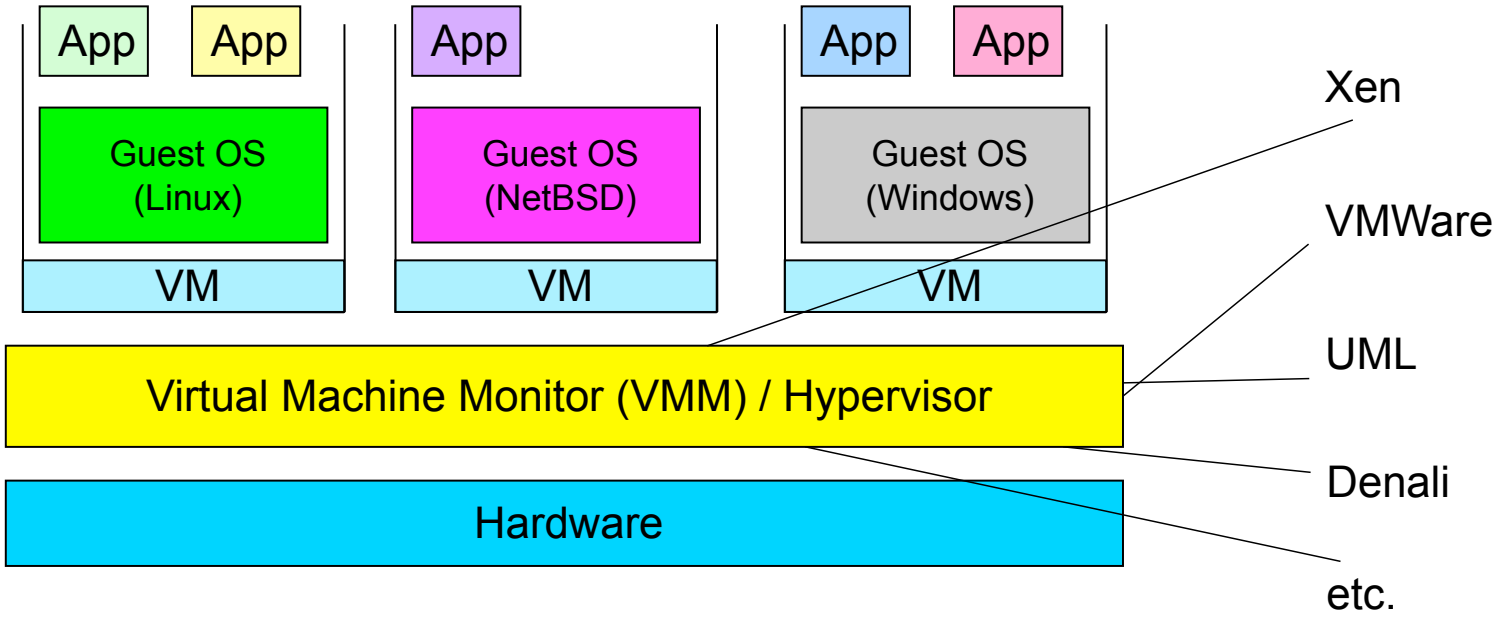

Performance: Para-virtualization (e.g. Xen) is very close to raw physical performance!

Windows Azure Components

	Windows Azure PaaS
Applications	Windows Azure Service Model
Runtimes	.NET 3.5/4, ASP .NET, PHP
Operating System	Windows Server 2008/R2-Compatible OS
Virtualization	Windows Azure Hypervisor
Server	Microsoft Blades
Database	SQL Azure
Storage	Windows Azure Storage (Blob, Queue, Table)
Networking	Windows Azure-Configured Networking

adopted from [3]

Datacenter Architecture

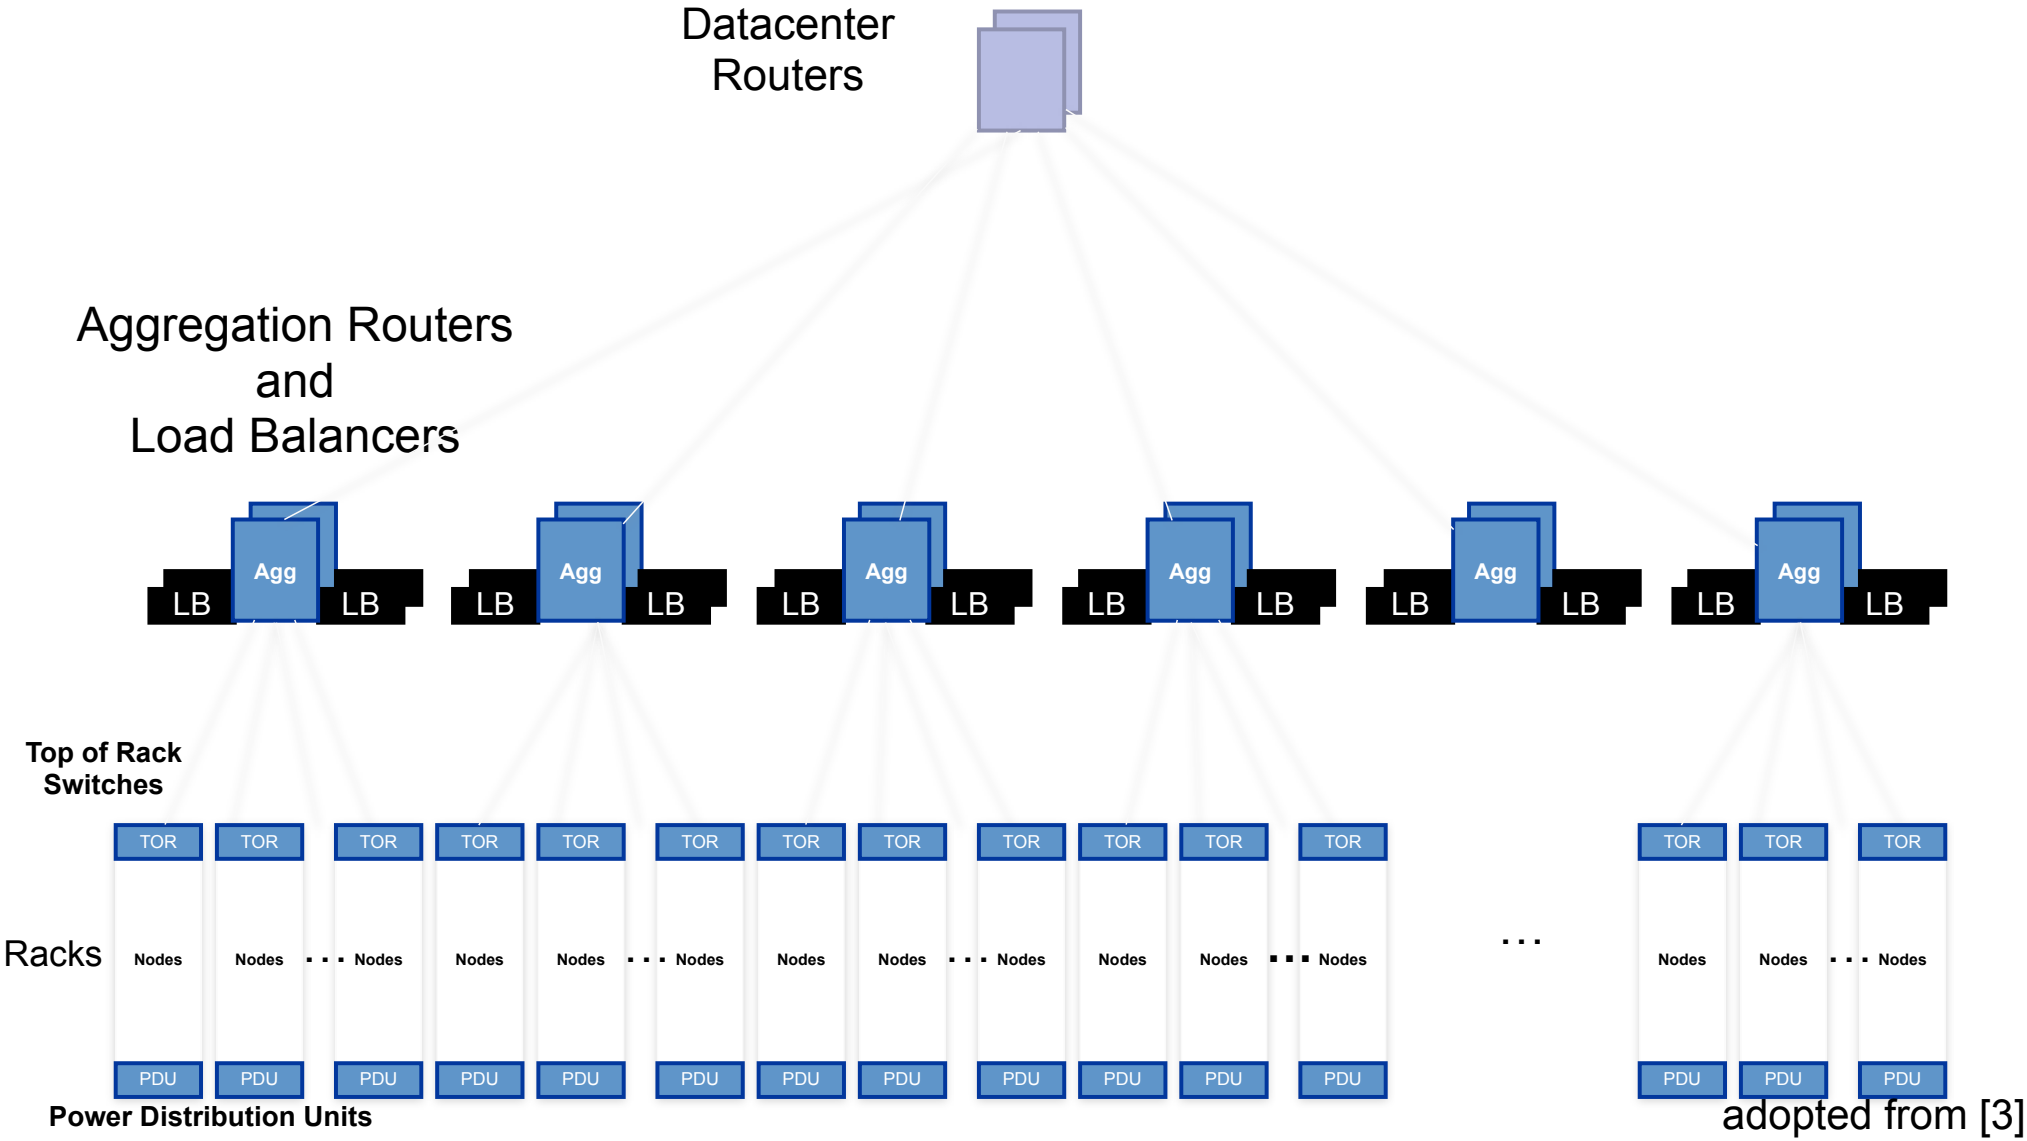

adopted from [3]

Windows Azure Datacenters

adopted from [3]

benefits of cloud computing

- enables companies and applications, which are system infrastructure dependent, to be infrastructure-less
  - instant software updates
  - unlimited storage capacity
  - robust against client's local hardware failures
  - clients can access their data and services from anywhere
  - easier collaboration over data in the cloud
- save in capital and operational investment
- clients can:
  - put their data on the platform instead of on their own desktop PCs and/or on their own servers
  - they can put their applications on the cloud and use the servers within the cloud to do processing and data manipulations etc.

disadvantages of cloud computing

- requires constant and fast Internet access
- can be slow
- features might be limited
- stored data can be lost
- inappropriate for some applications (e.g., HPC)
- applications have to be adapted to the cloud infrastructure and APIs
- security concerns

how public cloud security differs

adopted from [2]:
Charlie Kaufman “What’s different about security in a Public Cloud (Compared to a conventional data center),” keynote at Cloud Computing Security Workshop, October 2011.

What's Different?

- The stakes are higher
- The customers are less trusted…
  - Must be treated as hostile
- The customers' data must be protected from system operators
  - What's good practice within an enterprise is a contractual guarantee in a public cloud

What's the Same?

- Detecting and preventing intrusions
- Mitigating DDoS attacks
- Protecting services from one another
  - Including fair allocation of shared resources
- Keeping patches up to date
- Focus on minimizing the attack surface

Division of Responsibilities

- Protection of a service requires use of a variety of tools
  - Some can be used by the cloud provider
  - Some can be used by the customer
  - *Some can't be provided easily by either, and these require some workaround*

Generic Cloud Computing Engine

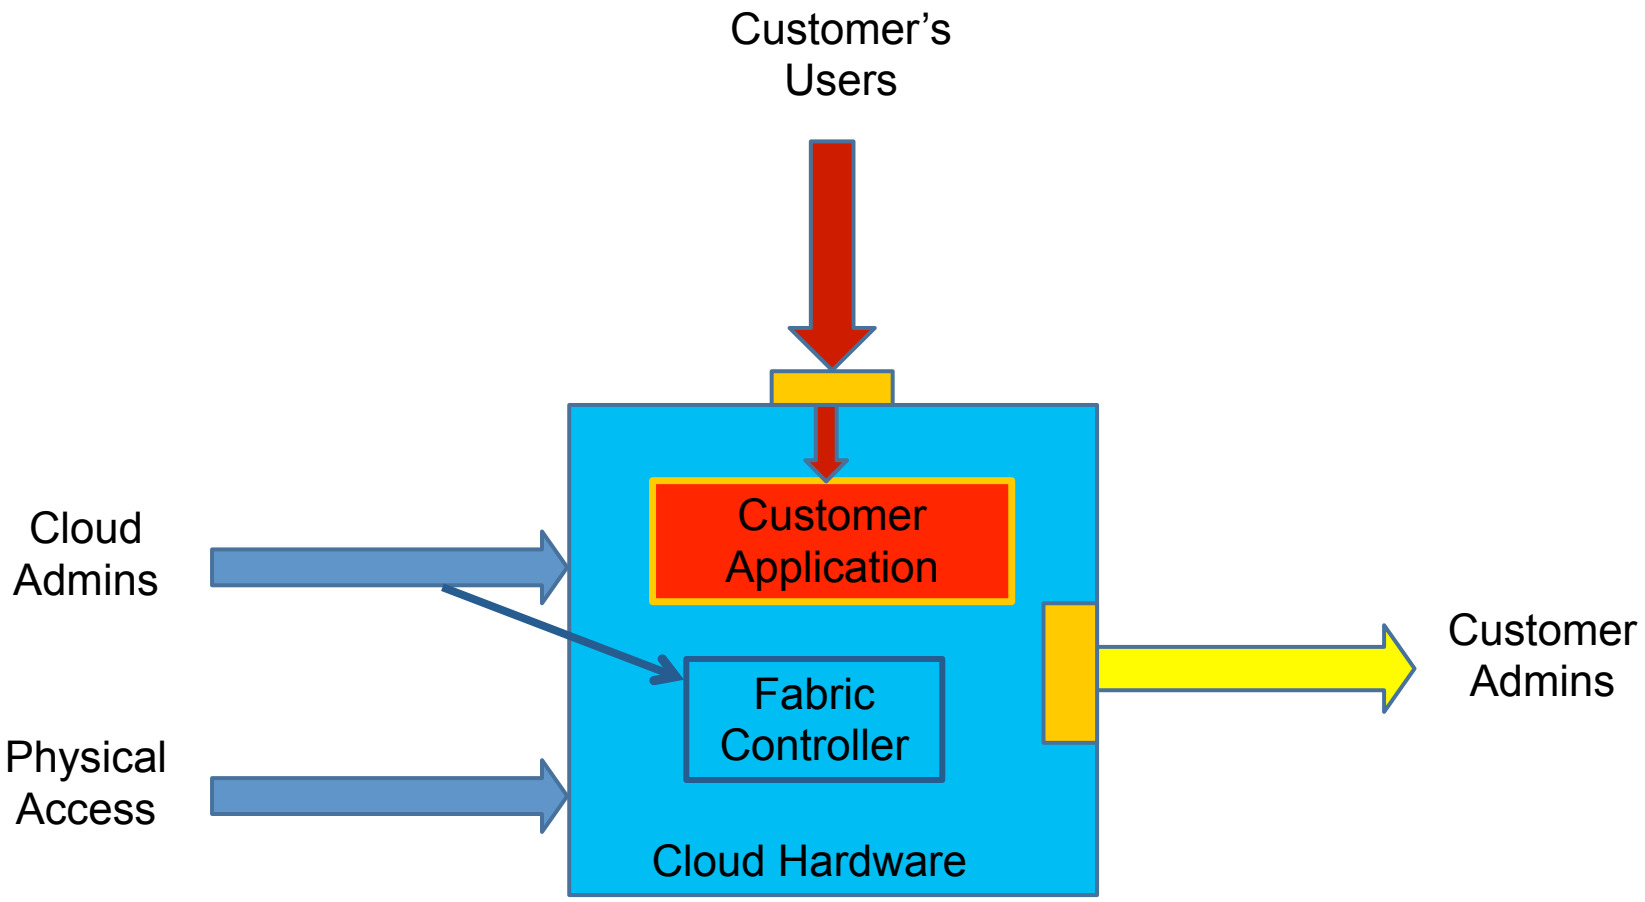

Generic Cloud Computing Engine

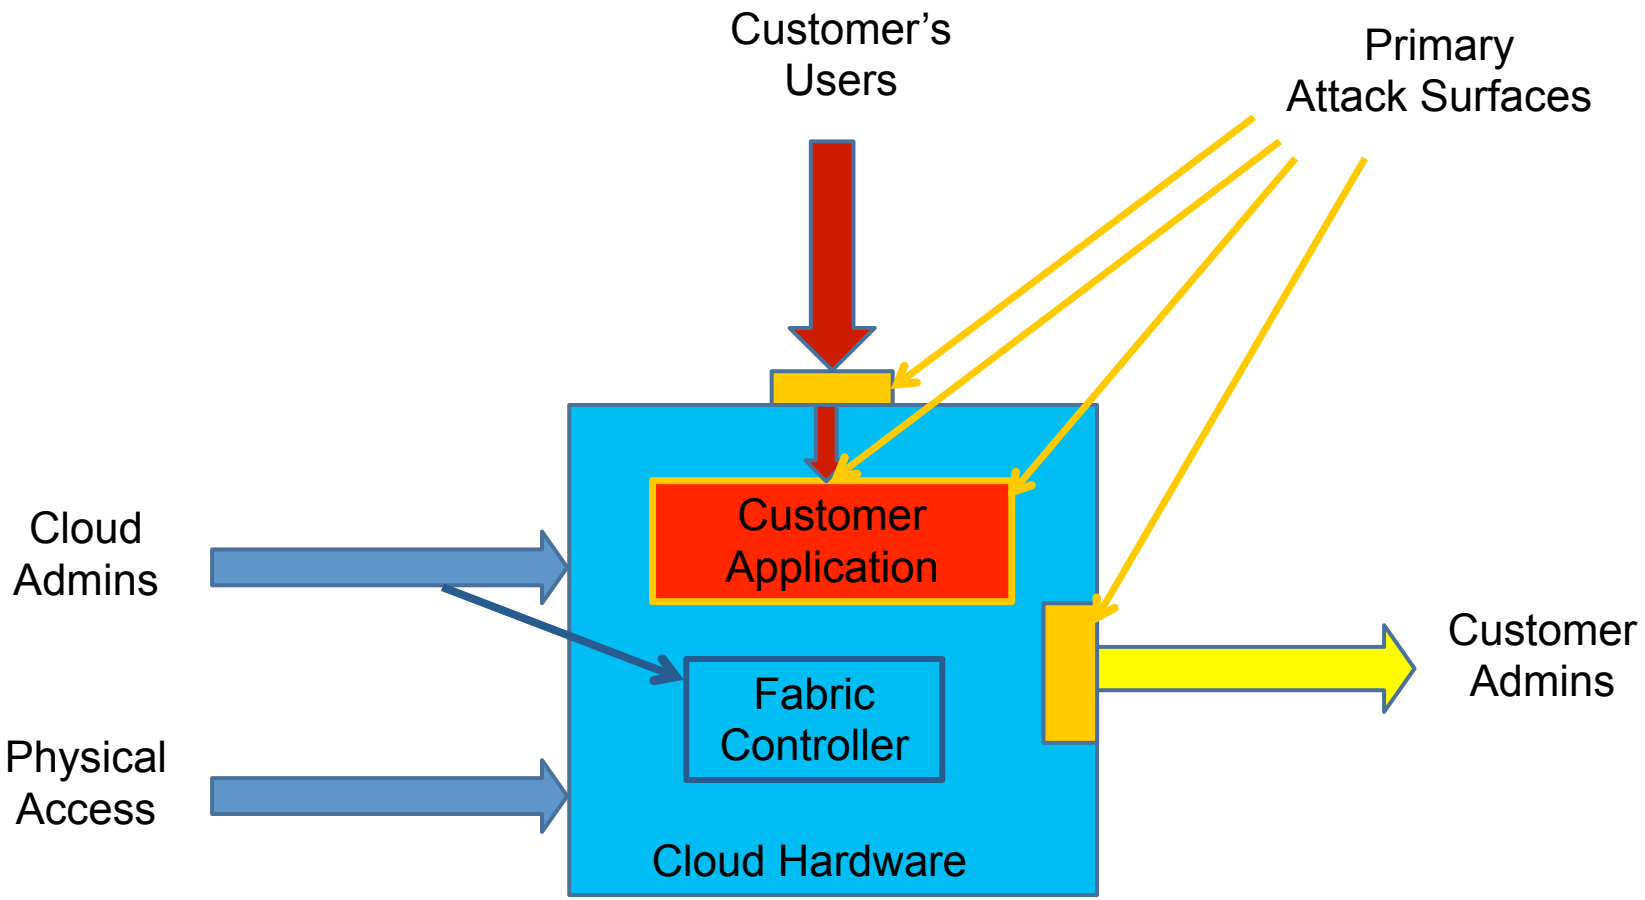

XML signature wrapping attack

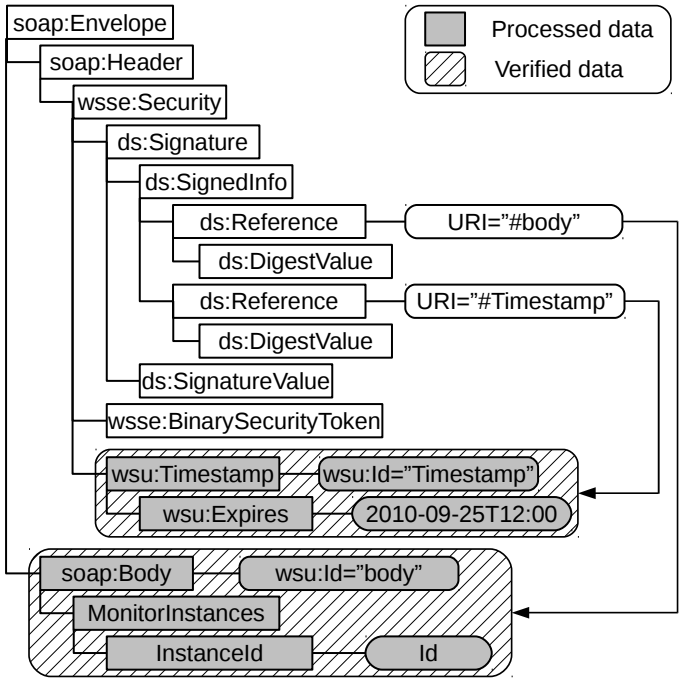

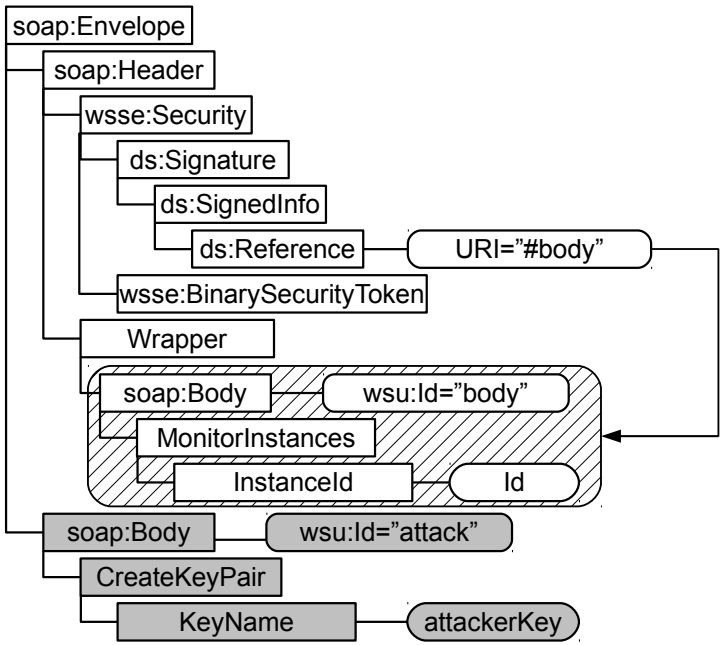

adopted from [4]

place of check to place of use (POCTPOU)

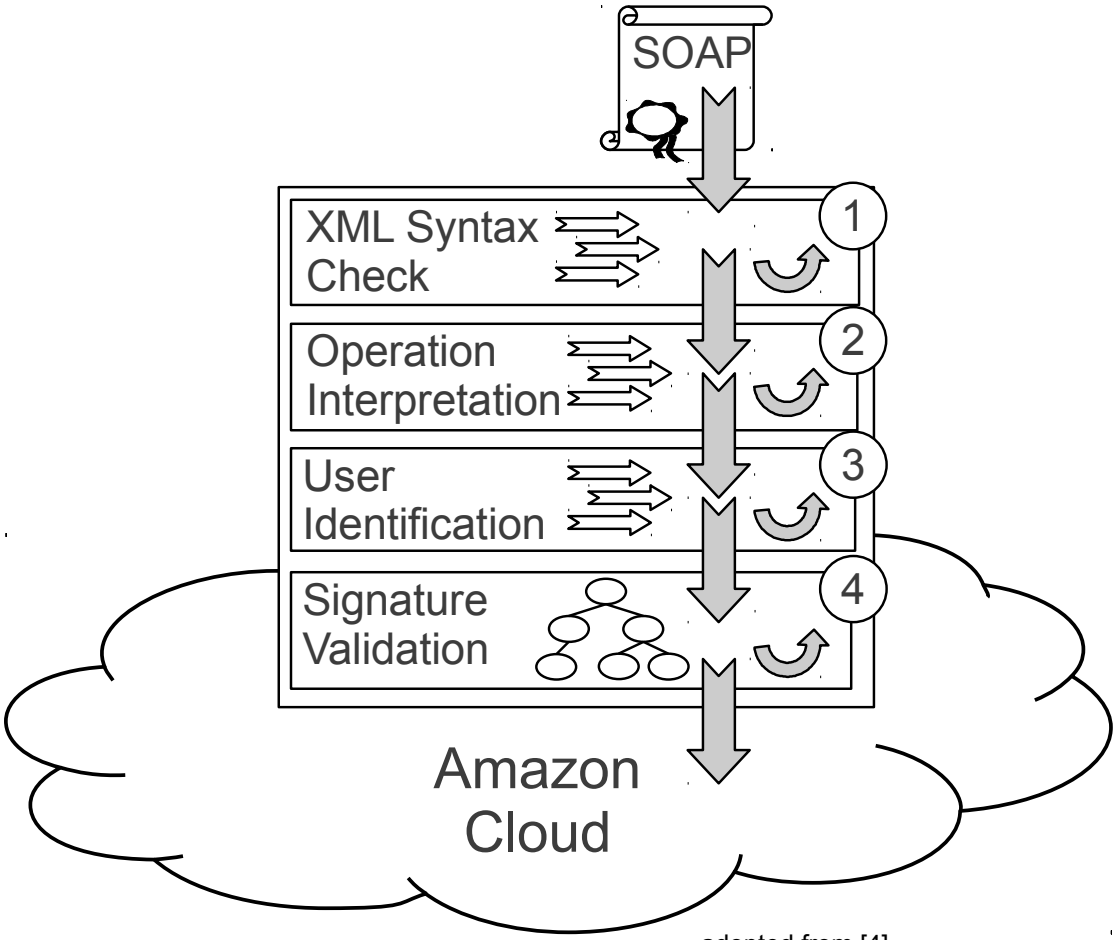

adopted from [4]

time of check to time of use (TOCTTOU)

Protecting the Infrastructure from Customer Admins

- Many systems delegate limited administrator privileges
- …but they typically don't assume the limited administrators are actually hostile
- In a public cloud, you must assume they are

Protecting the Infrastructure from Customer Applications

- Within a corporate data center, it is not unusual for some server to be compromised by some bug
- Designers therefore should assume that these applications might be hostile
- But most don't take the threat seriously; in a public cloud, we must
- If you mess up in your own data center, you're less likely to be sued

Helping Customers to protect themselves from their users

- Typical datacenters don't expose their servers to the full onslaught of the Internet
  - Datacenter firewalls
  - Intrusion detection hardware/software
  - DDoS mitigation systems
  - SSL accelerators
- Often these require considerable expertise to configure optimally

So those were the attacks we prepared for…

What did we actually see?

- Bots establishing accounts with stolen credit cards
  - A new challenge that requires some innovative thinking…

Protecting the Internet from our Customers

A Cloud provider acts as – among other things – an Internet Service Provider

- Provides greater anonymity than most ISPs
- Provides more bandwidth than most ISPs
- Rents out resources for a much shorter period of time

What kinds of behavior are acceptable?

Bad Behavior

- Acting as a rendezvous point for a bot army
- Impersonating another site in a phishing attack
- Sending out Spam!
- Posting malware for download
- Conducting DoS attacks (AaaS)
- Probing systems for vulnerabilities

The Internet has developed an immune system

- IP addresses that are the source of spam or malware get blacklisted
- IP addresses that are the source of DoS or probing attacks are blocked and reported to their owners for corrective actions
- If someone rents an IP address and a gigabit of bandwidth for 15 minutes, the reaction hurts the next tenant

How do you define bad behavior?

- How do you distinguish a spam engine from a mail agent relay distributing mail to a mailing list?
- How many failed DNS queries are allowed before it constitutes an exhaustive search through a namespace?
- What looks like an attack could be someone testing the security of their own system

How do you handle complaints?

- Forward them to the customer responsible?
- Forward customer contact information to the complainant?
- The complainant could be complaining as a form of DoS attack on the customer

credits

1. Mark Baker “An Introduction and Overview of Cloud Computing” ACET, University of Reading, 2009-05-19.
2. Charlie Kaufman “What’s different about security in a Public Cloud (Compared to a conventional data center),” keynote at CCSW ’11, October 2011.
3. “Introducing Windows Azure” Arnon Rotem Gal-Oz, Alon Fliess.
4. J. Somorovsky, M. Heiderich, M. Jensen, J. Schwenk, Nils Gruschka, and Luigi Lo Iacono, “All your clouds are belong to us: security analysis of cloud management interfaces,” In Proceedings of the 3rd ACM workshop on Cloud computing security workshop (CCSW ’11). ACM, New York, NY, USA, 3-14.