Cloud Network Security: Challenges and Opportunities

Charlie Kaufman

Windows Azure Security Architect
Microsoft

Definition of Cloud

Definition of Cloud

Cloud (n): vapor in the sky

Definition of Cloud

Cloud (n): vapor in the sky

Cloud (v): to make unclear or confused

Definition of Cloud

Cloud (n): vapor in the sky

Cloud (v): to make unclear or confused

Definition of Azure

Azure (adj): deep blue, like the color of the sky on a cloudless day

Definitions of Cloud Computing

- Different people mean different things
  - Serve clients on the Internet
  - Massive highly available services (e.g. Facebook, Google, Hotmail)
  - Run on Virtual Machines (over hypervisors)
  - Remotely managed
  - Datacenter Operator separate from Service Operator

Definitions of Cloud Computing

- What I mean is services that are dynamically created, deleted, and scaled and where all service management is remote
- Usually implemented with large homogeneous data centers with virtual servers and virtual networks

How is Cloud Security Different?

- For a private cloud, not a lot:
  - Physical security is separately managed and better
    - Because only hardware repair people ever need to go there
  - All the old threats are still there
  - The ubiquitous control structure is an attractive target
  - Homogeneity can lead to catastrophic failures
  - Less room for add-on security "appliances" like firewalls, intrusion detection systems, etc.

How is Cloud Security Different?

- In a public cloud, security becomes much more important
  - More like commercial time sharing systems from the ‘70s
  - The cloud operators and customer administrators are mutually distrustful
  - The attackers may be on the inside of your network
  - Cloud administrators may not be authorized to see customer data – this makes diagnosis tricky
  - Get security wrong and you don’t have a viable business
  - Availability is key – Denial of Service cannot be an afterthought

What does the hardware look like?

- Blades, Switches, and Racks
- A Blade has some CPUs, memory, disks, and one or two network ports
- A Switch has a large number of network ports
- A rack holds a power supply, lots of blades, and a few switches

What does the hardware look like?

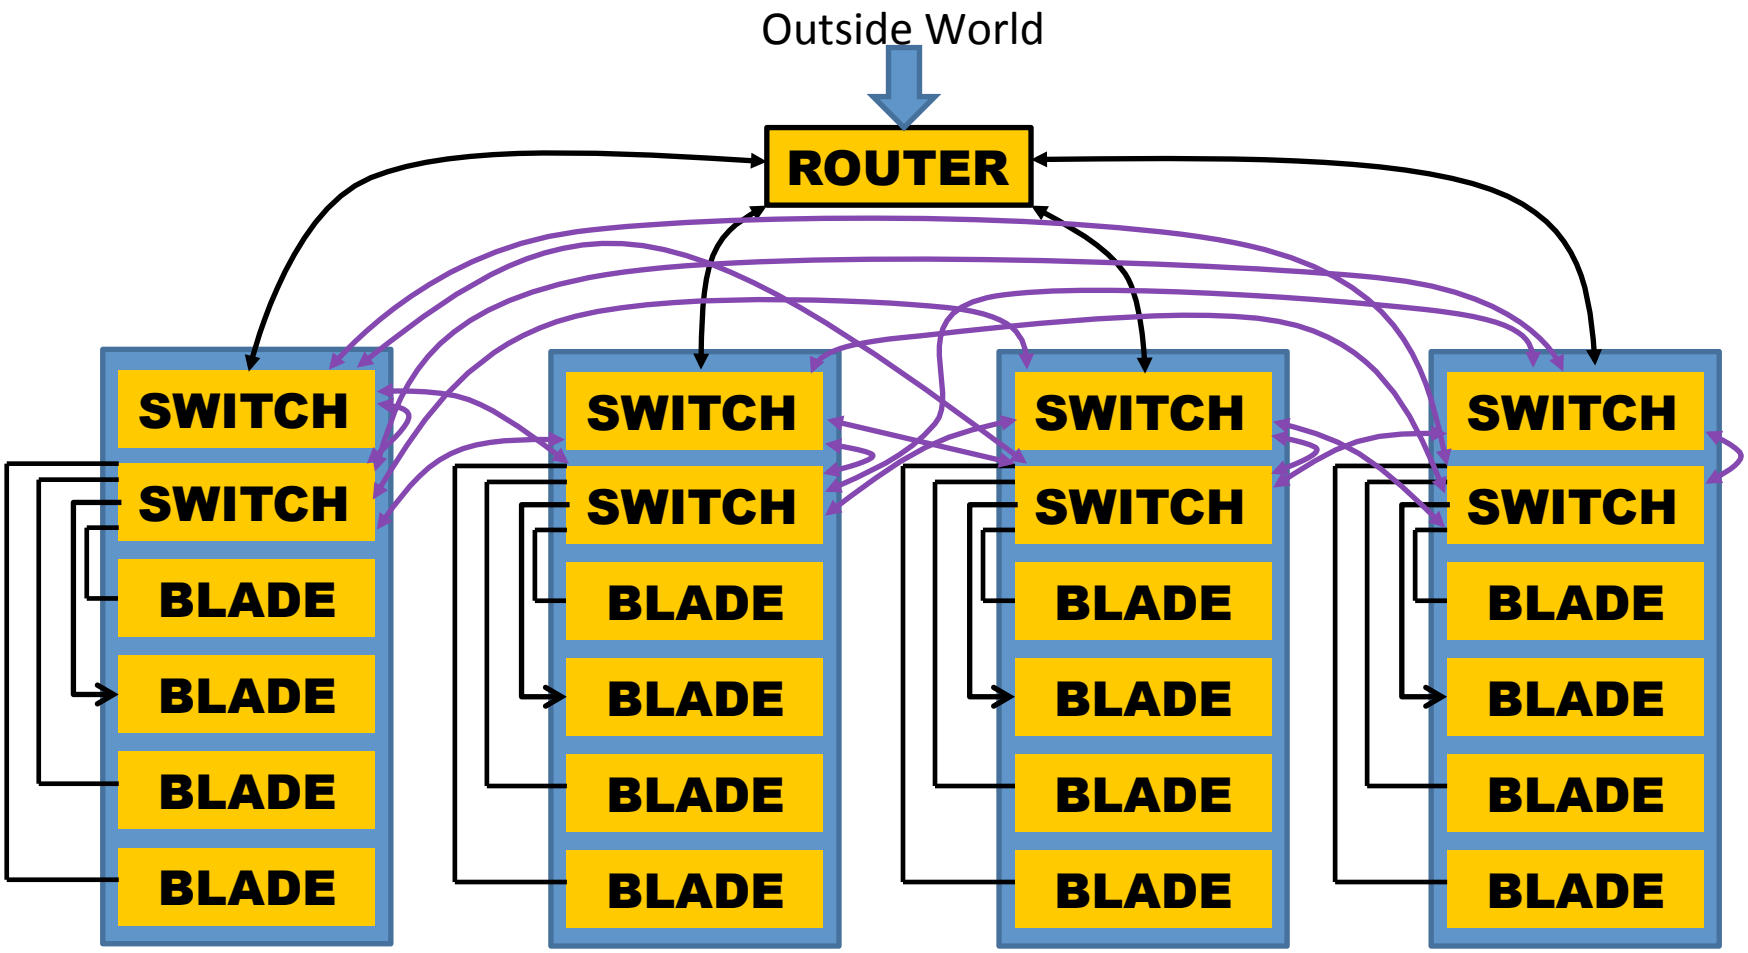

Typical Cluster: 10K Cores, 20TB memory, 10PB Disk Space
Think of this as your typical server

Virtualized Network

- Each customer deployment sees its peer VMs and no others
  - Might see the Internet, or an internal network, or both
- VMs can move from place to place but their addresses don’t change

What does someone typically use a Public Cloud for?

A Deployment “On the Internet”

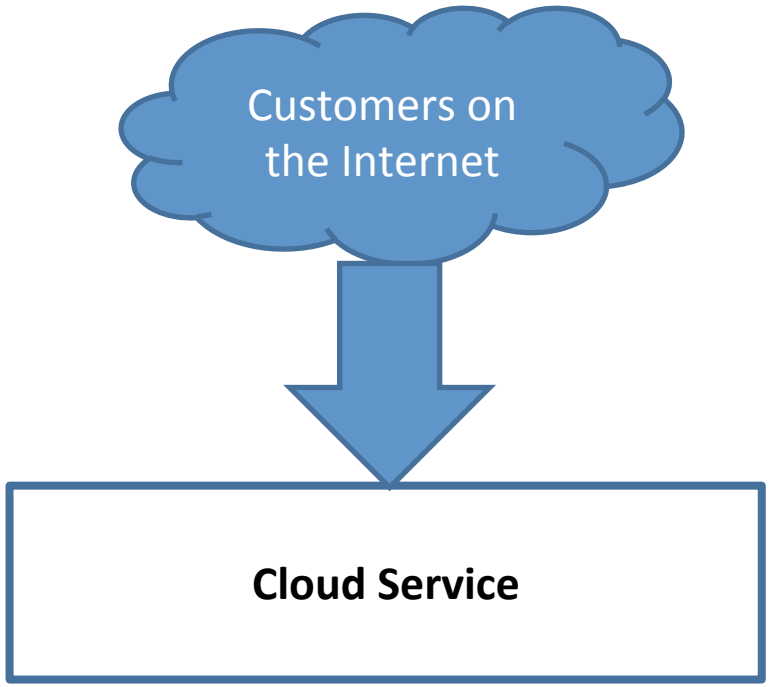

Adding Resources to an Intranet

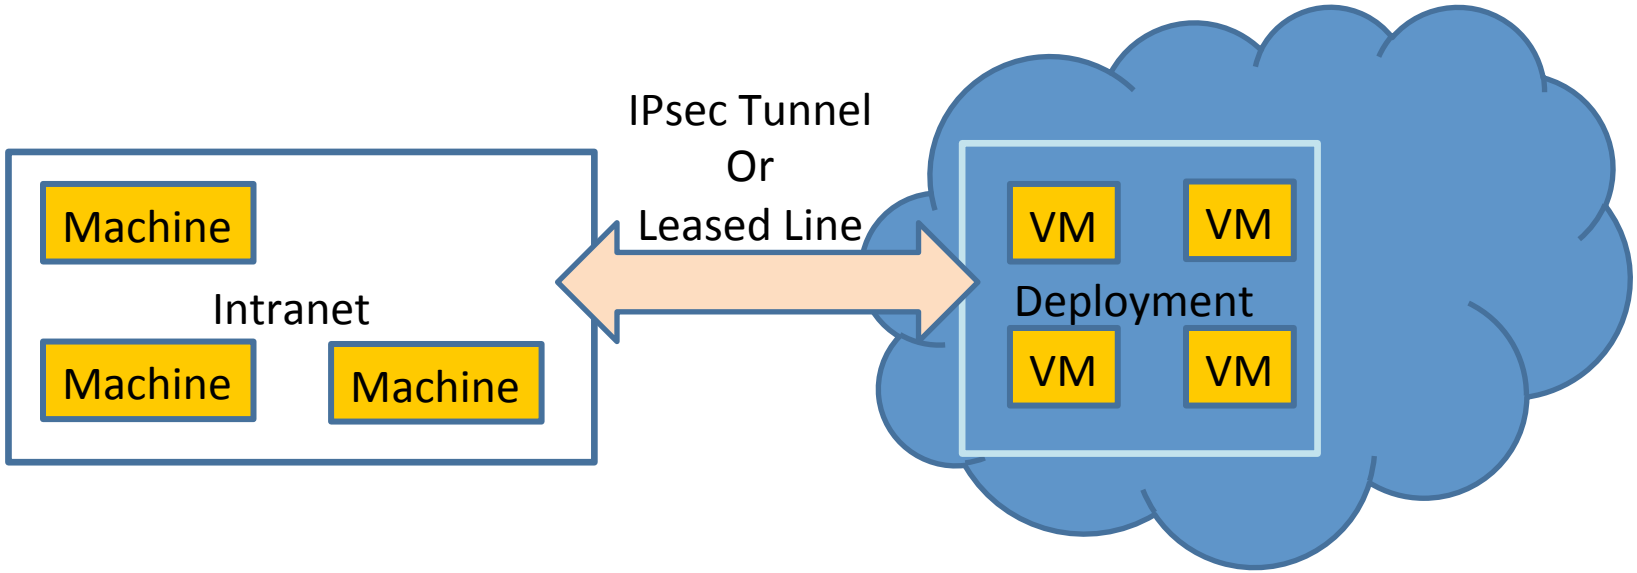

Using Cloud Resources as a DMZ

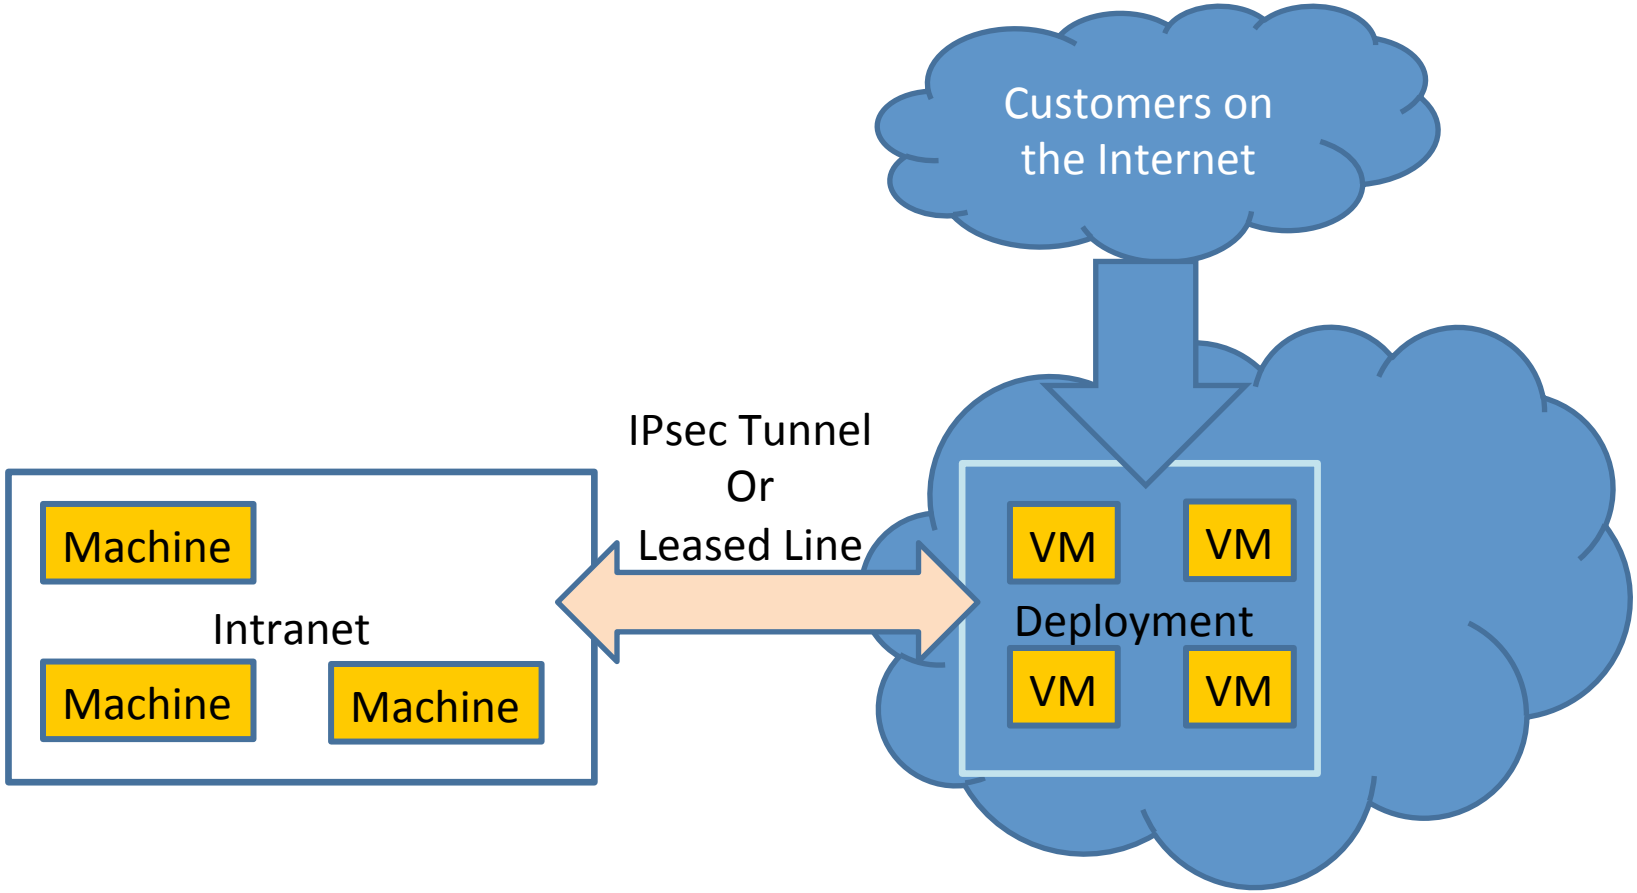

Think of the Internal Network as a Backplane

- Must be secure without encryption
- Eavesdropping and node impersonation must be prevented
- Two ways to get there:
  - No untrusted code connected directly to network
    - Hypervisors or trusted OSes block unauthorized traffic
  - Switches enforce isolation rules
    - VLANs per customer deployment
    - IP level filtering in switches

Autoconfiguration vs. security

- Traditional Ethernet: Learn from location of seeing a packet with a source address; flood if unknown
- VLAN isolation enforced by switches
  - Except sometimes endnode has to (hypervisor with multiple VLANs possible), but switch can check that it's an allowed value
- Clouds: Want to have a fabric manager that knows where everything is and configures switches and hypervisors

Biggest Decision in the Design of a Cloud OS

- How to isolate customer applications from one another
  - Commonly with a hypervisor, but there are other options
  - At machine boundaries if you can trust the network to enforce isolation
  - At process boundaries if you can trust the OS to enforce isolation
  - With type-safe code within a process

What would we need from the network?

- Ideally, network controls connectivity with VLANs or IP ACLs
- Software Defined Networking is ideal for this sort of rules
- At a minimum: prevent unauthorized parties from forging a source IP address or eavesdropping on traffic not addressed to them

Isolating Customer Data

- Typically want to interleave data from multiple customers on shared disks
  - Virtual disks
  - Virtual File Systems
  - Virtual databases
- Bugs could leak customer data to other customers
- A per-customer cryptographic keys minimizes the chances of accidental disclosure
- Bugs could still corrupt data, which is usually worse

Generic Cloud Computing Engine

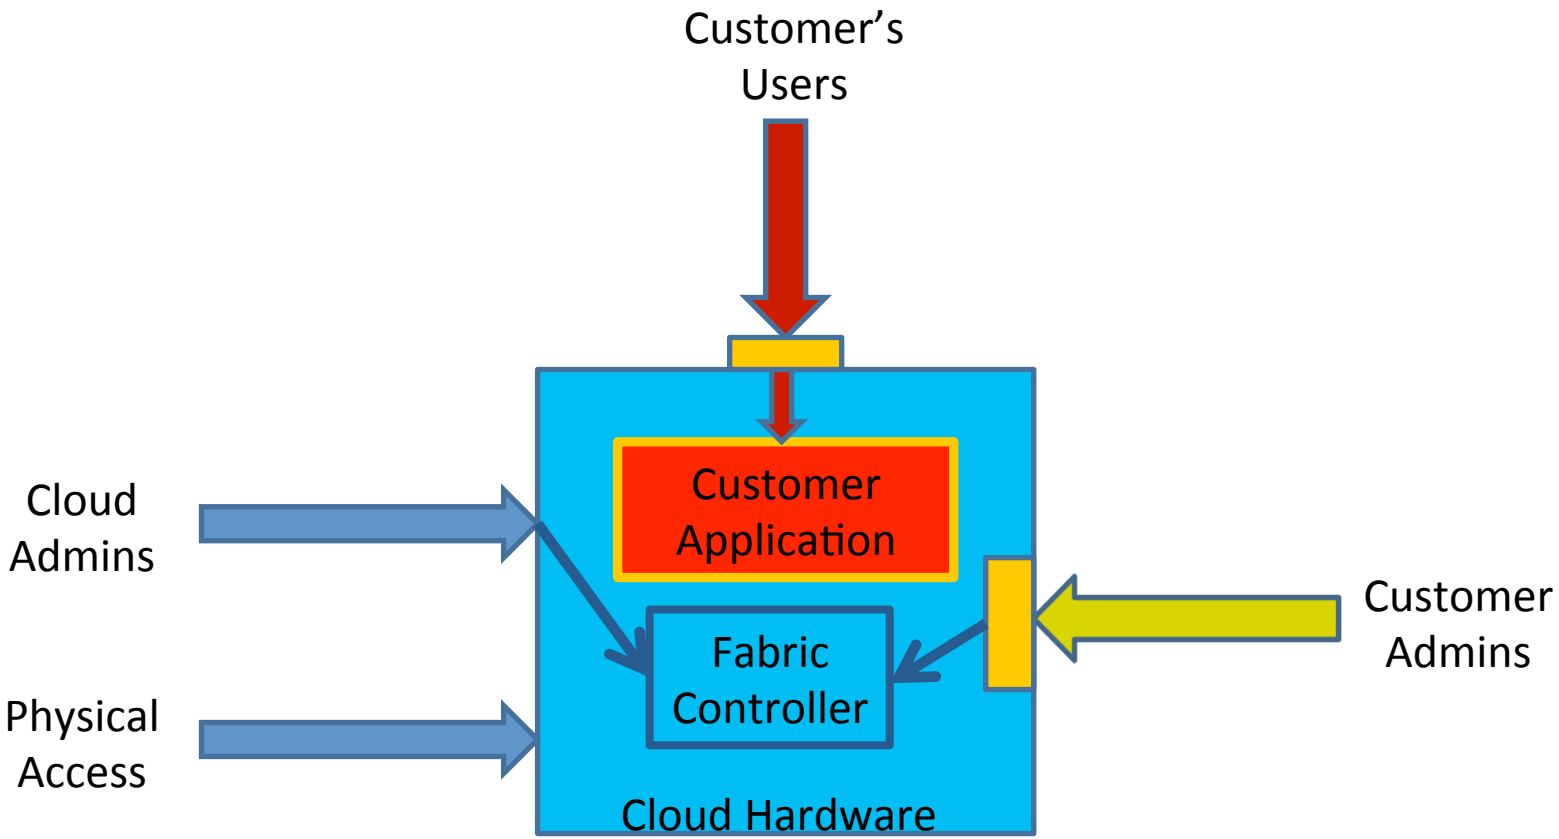

Generic Cloud Computing Engine

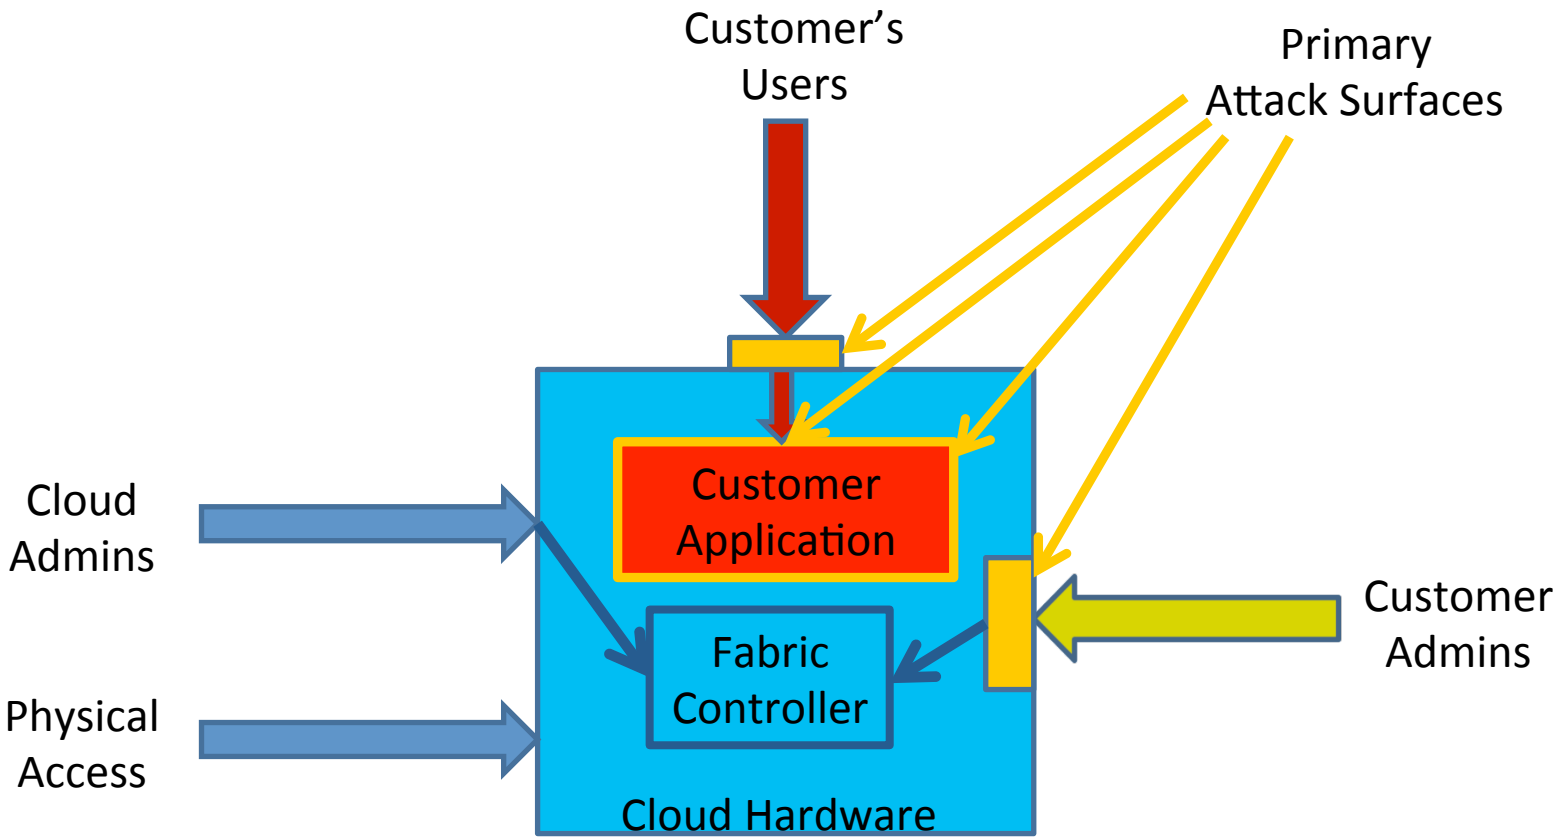

Protecting the Infrastructure from Customer Administrators

- 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
- Have to worry about Denial of Service and scheduling fairness

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
  - DoS mitigation systems
  - SSL accelerators
- Often these require considerable expertise to configure optimally

Challenges with Load Balancing SSL

- SSL was designed so that repeated connections from the same client to the same server efficiently
- Load balancers typically send connections from the same client to multiple servers
- Unless servers coordinate carefully, this results in very poor SSL performance

Solutions for Load Balancing SSL

- Offload SSL processing to high capacity front end processors
- Have servers share a session cache
- RFC5077 – TLS protocol enhancement
  - (not yet widely implemented)

Intrusion Detection

- It is mostly the responsibility of individual services to detect assaults
  - Password guessing
  - SQL injection
  - Mal-formed requests
- Some things a front end can do better
  - Keep up to date with attack patterns
  - Notice similar attacks on many services
  - Difficult if SSL processing is on server

Dealing with DoS

- Denial of Service attacks are the last thing security people think about but the first attack they will experience
- Lots of categories of attacks calling for different responses
- Case 1: Flash Mob… quickly scale up the capacity of the web site

Dealing with DoS

- Case 2: More bandwidth in fake traffic than a site can handle
  - Filter the traffic upstream, before it reaches a bottleneck
  - May be able to identify good / bad traffic
  - If all else fails, protect customers not under attack by limiting bandwidth

Dealing with DoS

- Case 3: Sometimes a site has trouble with a low bandwidth attack and a front end can help
  - Sites could prioritize based on IP address or close largely idle connections, but most don't
  - A well designed reverse-proxy front end may be able to insulate a server
    - SSL processing offload
    - Assemble and validate a request before forwarding it to the service

So those were the attacks we prepared for…

What did we actually see?

So those were the attacks we prepared for…

What did we actually see?

1. Bots establishing accounts with stolen credit cards

So those were the attacks we prepared for…

What did we actually see?

1. Bots establishing accounts with stolen credit cards
2. Using the stolen accounts to do bad things on the Internet

So those were the attacks we prepared for…

What did we actually see?

1. Bots establishing accounts with stolen credit cards
2. Using the stolen accounts to do bad things on the Internet
   - DoS attacks
   - Password Guessing / SQL Injection
   - Sending Spam
   - Spreading malware or copyrighted material

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

Dealing with DoS

- Case 4: Attack is Outbound
  - How to distinguish good traffic from bad?
  - Sometimes it's easy:
    - Large numbers of failed connection attempts is either a port scan or a DoS attack
    - Lots of short connections to ports that conventionally accept username and password credentials is highly suspicious

Challenges with Copyrights

- When presented with a take-down notice, a hosting service has a deadline to respond or be held liable
- We won't divulge the identities of our customers without a court order
- Is a cloud provider more like a hosting service or more like an ISP?
- How do we maximally protect our customers while staying within the law?

When the Customer Doesn't Trust the Cloud Operator

- A customer may worry that the cloud operator will peek at the customer's data
  - Or may worry that the cloud operator's government will insist on doing so
- Must defend against a rogue employee with dual controls, extensive auditing, and "least privilege" roles
- But is there anything a customer can do?

When the Customer Doesn't Trust the Cloud Operator

- Can use the cloud only for data storage (perhaps only backup) and encrypt all data with keys not placed in the cloud

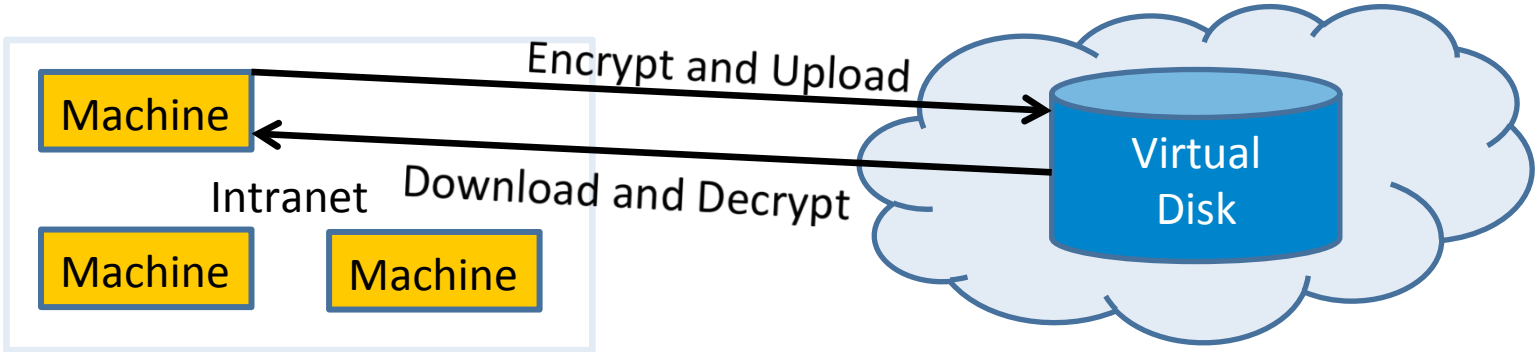

When the Customer Doesn't Trust the Cloud Operator

- Encrypt the data, but not the meta-data
- Upload my photos encrypted, but don't encrypt textual descriptions
- Use text search to find the photos of interest, then decrypt after downloading
- (Doesn't help if the textual descriptions are secret)
- Research efforts searching encrypted data

When the Customer Doesn't Trust the Cloud Operator

- Living dangerously… only keep the key in the cloud while active (for instant revocation)

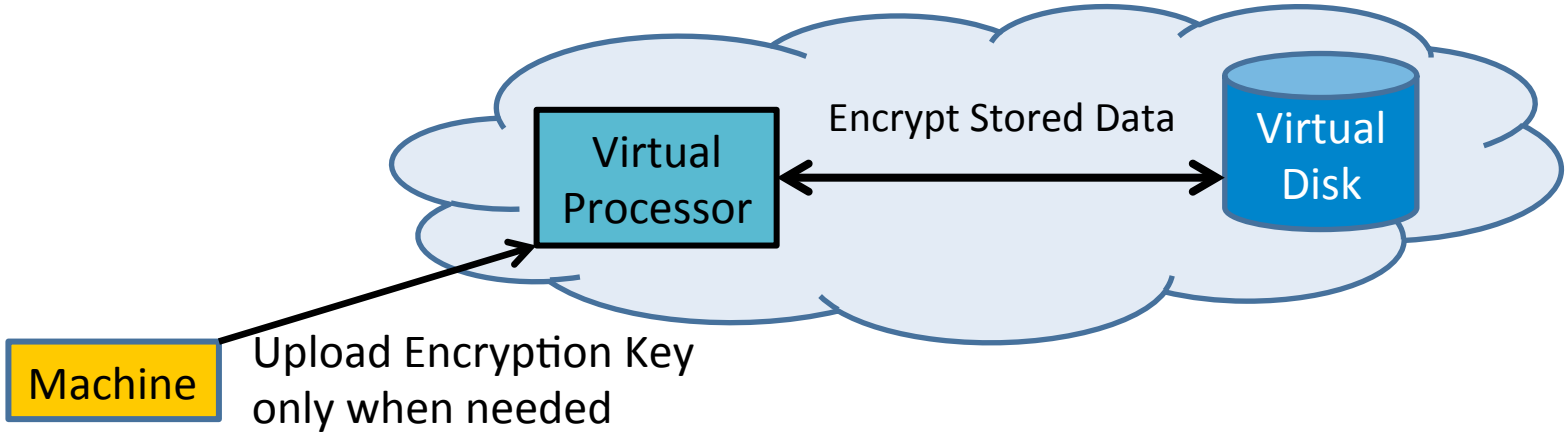

A Pipe Dream: Homomorphic Encryption

- With some modes of RSA,

  Encrypt(A) * Encrypt(B) = Encrypt(A*B)

- Means you can upload encrypted values to the cloud, have a cloud VM multiply arbitrary subsets together and return the results for decryption, and the cloud VM never has access to the plaintext data

A Pipe Dream: Homomorphic Encryption

- Suppose you could invent a cryptographic algorithm and operations ★ ✚ & such that

  Encrypt(A) ★ Encrypt(B) = Encrypt(A*B)

  Encrypt(A) ✚ Encrypt(B) = Encrypt(A+B)

  Encrypt(A) & Encrypt(B) = Encrypt(A&B)

- Then you could do arbitrary calculations in the cloud without exposing any plaintext data

A Pipe Dream: Homomorphic Encryption

- Cryptographic functions and operators like this *may* be found
  - Cryptographers have found functions almost like this, but not quite
- But… the performance considerations make it unlikely to ever be practical

Questions?

Backup

Abstract

Engineering a network for a large public cloud computing facility presents some unique challenges. Addressable entities like virtual machines need to migrate from place to place while keeping their network connections alive, which would tend to require a highly dynamic routing algorithm with fast convergence. But a primary goal of cloud computing is to reduce costs by pushing all components to their limits, which means maximizing bandwidth, avoiding bottlenecks, and load splitting across parallel paths - all of which argues for a highly tuned highly static configuration. Security challenges take on new forms in a public cloud because the attackers may be inside your network and you can't profile what are "normal" usage patterns because your customers don't have to tell you what they are doing. Even diagnosing problems can be made difficult because you have to respect the privacy of your customers even as you try to determine whether they are trying to intentionally harm your network.

Fortunately, with these challenges come new tools. Data centers are large and homogeneous, without the need (or ability) to rewire things to deal with some new hardware rolling in. Management is largely automated, avoiding the problems of administrators working at cross purposes without knowing of one another's existence. Perhaps most importantly, no untrusted software is directly on the network. Because all network access is funneled through network monitors beyond the reach of even privileged users on a VM, many types of mischief can be blocked, monitored, profiled, and rate limited. This talk will describe some early experiences in this brave new world.

What do I mean by the Cloud?

- Everyone you talk to will mean something different
- I mean large homogeneous (hardware) data centers where management software creates diverse virtual environments where virtual machines, virtual networks, and virtual storage hides the hardware from developers
- No one knows or cares where the hardware is
- Easy scaling over a vast range of sizes

Hardware Vendors vs. Cloud Operators

- Some fundamental miscommunications
- Hardware vendors want to know what features will command a premium price
- Cloud Operators are trying to wring every dime out of hardware cost
  - Compensate for unreliability with redundancy
  - Compensate for missing features with virtualization
- Commodity hardware is becoming feature-rich

What do Physical Layer Addresses look like?

- Could be:
  - Ethernet
  - IPv4
  - IPv6
  - Infiniband
  - …
- We don't care… whatever is cheapest and performs best
- Virtual Addresses tend to be IPv4… it's what is most familiar to customers
  - It may be IPv6 in the future

The Network is a Backplane

- Nothing untrusted is connected to it
- Hypervisors on Blades virtualize network access
- Data transferred may not be cryptographically protected yet still must be secure
- Network infrastructure components (routers and switches) must be secure
  - Favors limiting programmability

What do we need from a Backplane?

- Prevent source spoofing and eavesdropping
  - As a defense in depth
  - To permit unshared blades to be unvirtualized
  - DHCP Option 82 a good basis for protection
- Denial of Service and Traffic Shaping a BIG Deal!
  - Difficult to predict what statistics and filtering will be needed
  - Sampling of dropped packets to identify congestion sources

DHCP Option 82

- First switch adds a field to DHCP request indicating physical port of connection
- DHCP server assigns IP addresses based on physical ports
- Switches notice and remember the IP addresses assigned to node on a port
- Nodes cannot send packets with IP source addresses not assigned to them by DHCP
- Nodes cannot receive packets with IP destination addresses assigned to them by DHCP

Alternative Network Isolation Techniques

- Trusted infrastructure places VMs on nodes, and knows where everything is
- It would explicitly tell switches what IP addresses and Ethernet addresses are associated with each physical port
- Switches would need to accept a high update rate
- Autoconfiguration is your enemy