

Cloud Services



The Cloud: Threat or Menace?

- Everyone loves “the Cloud”—it’s the key to secure, scalable computing, omnipresent data, and more
- Everyone hates the Cloud—it’s a threat to security and privacy. “The Cloud is other people’s computers.”
- Is it both? Neither?

The Cloud

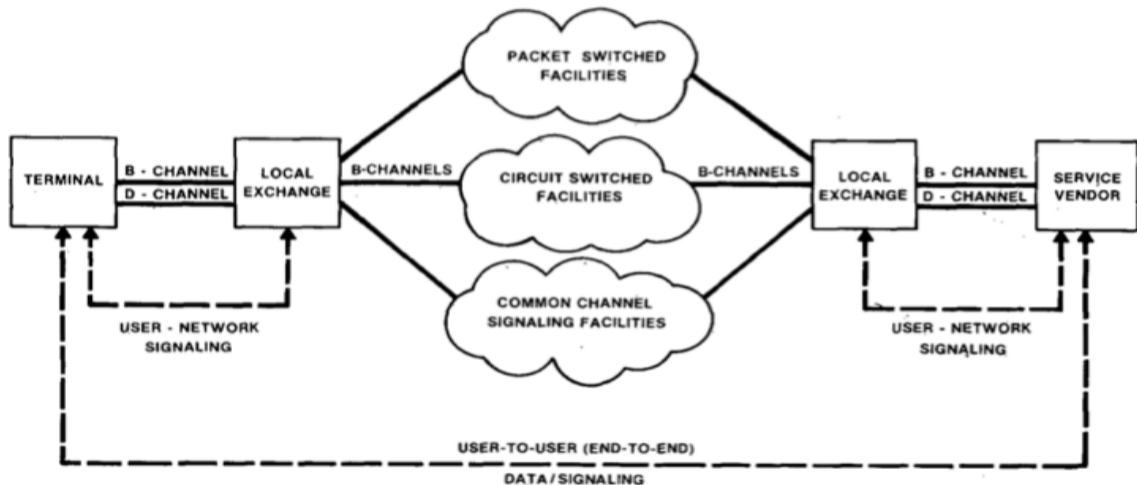


Fig. 7. The three-cloud network.

What is “The Cloud”?

- A cloud is a traditional way to represent a network
- This “three-cloud network” picture is from 1982
- But: today “cloud” refers to **computing services** provided via the **Internet** by an **outside party**
- (The modern usage seems to date to 1996)

“Via the Internet”

- The service is not provided on-premises
- An Internet link is necessary
- This link might itself be a security risk, especially for availability

“Outside Party”

- By definition, cloud services are provided by an outside party
- *Not* the same as a company's own remote computing facility
- Organizations can have a “private cloud”, but the security issues are very different

Types of Service

- Storage
- Computing
- Applications
- Virtual machines
- More

- Disk space in a remote location
- Easily shared (and outside the corporate firewall)
- Often replicated for reliability
 - Replicas can be on different power grids, earthquake zones, countries, continents, plague areas, etc.
 - Data can be moved to be closer to users
- Expandable
- Someone else worries about space, backups, security, etc.
- Examples: Dropbox, Google Drive, Carbonite, Amazon S3

- Rent cycles as you need them
- Pay only for what you use
- Often used in conjunction with provider's cloud storage service
- Examples: Amazon EC2, Microsoft Azure, Google Cloud
- (Dropbox is a cloud service that uses a different provider's cloud storage)

- Provider runs particular applications for clients
- Common types: web sites, email services
- Less common: shared word processing, payroll
- Examples: Gmail and Google Docs, Microsoft Outlook and Office 365, Dreamhost (web hosting)

Playing an Active Part: Google Docs

- Someone, using a web browser, creates a document
- Others who have proper authorization (or sometimes just a URL) can edit the document via their own browsers
- Changes made by one user show up in realtime for all other users

- Run many independent “computers” and operating systems on one box
- Each VM can be leased to a different customer
- Net effect: many computers that consume the space and power requirements of a single computer

There's Not One Cloud

- Each type of service can carry its own threats
- They often have to be analyzed separately

Location of Cloud Servers

- Responsiveness and effective bandwidth to a server is limited by distance
- The problem is the speed of light—it's too slow!
- Example: New York is about 4100 km from San Francisco—round-trip time is at least 40 ms (the speed of light in fiber is $\frac{2}{3}c$), plus router processing and queueing time
- 100 ms is perceptible
- It takes a minimum of 250 ms to set up a TLS connection from Washington to Paris, and twice that to New Delhi
- TCP bandwidth is related to distance, too
- Consequently, large cloud providers have server complexes all over the world
- (Also: located where space, power, and cooling are cheap)

Is Location a Risk Factor?

- What are the laws in the jurisdictions in question?
- What is the physical security in such places?
- Are the intracloud links encrypted?

- In the 1960s and 1970s, computers were big and expensive, and there was little commercial software
- Consequently, there were *time-sharing service bureaus*—dial-up computers
- They often offered specialized, canned applications for particular uses, e.g., payroll processing
- In other words: the Cloud, but via dial-up modem instead of the Internet

Is the Cloud Secure?

- What are we protecting?
- Against whom?
- And: is the Cloud more or less secure than in-house solutions?

- Using cloud services is not all or nothing
- What are you using the Cloud for?
- What are the indirect dependencies, e.g., resources accessible to someone who controls your email account?

The Enemies

- All of your usual enemies are still present
- Also: the Cloud provider
- Also: rogue employees of the Cloud provider
- Also: other customers of your provider
- Also: cloud-specific mistakes that you make

Common Threats

- A high percentage of system penetrations are due to buggy code, and in particular to bugs for which patches exist but have not yet been installed
- Who installs those patches on cloud systems?
- Often, cloud providers are *better* at patching than end-sites are
- For them, system administration *is* their job
- Cloud systems can thus be *more* secure
- But. . .

Patches versus Applications

- All too often, patches can break applications
- Patching quickly, then, has different properties for different cloud services
- For storage and applications, it's probably safe—*if* the provider has properly tested their own code
- For computing, where you're running your code on someone else's computers, it can be more problematic
- For virtual machines, you may be on your own even to install the patch

Patching in a VM Environment

- Never patch a running server
- Instead, patch a new VM and see if it works
- If it does, swap it in and shut down the old one
- If it seems to be failing, shut it down and revert
- (Or use checkpoint/revert functions of your VM platform, though you still need to test first)
- (Note that this can be done just as well on your own VM platform)

- What forms of authentication does your provider support?
- All of the usual issues apply—but you may not get to make that decision
- More precisely, if particular forms of authentication are important to your security, that has to be one of your selection criteria in picking a cloud provider

Cloud Storage and Sharing

- Who can access which cloud-resident files?
- Not authentication, *authorization*
- Most cloud storage systems permit sharing—but who can share, to whom, and how easily?

Sharing is Better with the Cloud

- Suppose you want to share some files with a joint venture partner
- You don't get to control which employees there are authorized
- How do you manage the authentication and authorization processes?
- Cloud-based sharing systems are already set up for that (or should be...)
- Otherwise, you'd have to set up a login system, authenticators, etc., and then manage it

- It turns out to be very easy to accidentally overshare on some platforms
- A Google search for "open aws bucket" get lots of hits
- How do you avoid making that mistake?
- Equally important: how do you monitor your own cloud storage to see if someone in your organization has accidentally done it anyway?

Employee Sharing

- Some organizations have sensitive internal data that they don't want employees taking home
- (That might even be a legal requirement)
- What if the employee uses a personal Dropbox account to export some data, for home access?
- Common solution: block Dropbox (and other cloud storage services) at the organizational firewall (though that's not hard for a determined miscreant to bypass)

Encrypting Cloud Data

- You can encrypt shared data—but who has the keys?
- The cloud provider—protects against other customers, and (maybe) some employees
- (Dropbox encrypts your files to protect them from employees of the underlying cloud storage provider)
- You, on the cloud machines—but a sysadmin there can probably get at your executing code
- Client-side, on your own computers—protects against the provider, but means you can't do computation in the cloud

- Probably the most common cloud service
- Many organizations (including, of course, Columbia) outsource their email to Google or others
- Why? Running a mail system is *hard*
- Issues: availability, security, spam filtering, more
- Many of these are easier at scale—and Microsoft and Google, two of the biggest, *really* understand scale
- But: is this secure?

Wrong Question!

- Organizations *need* email
- The right question is not “is cloud email secure?”, it’s “is cloud email more or less secure than doing it myself?”
- Almost certainly, Microsoft and Google can do it better than you can
- But: is the cloud provider your enemy?
- Specifically, what are their terms and conditions?

Cloud Email and Privacy

- Google, of course, scans email in free accounts to help target ads
- Microsoft relies on demographic data
- Most cloud email providers do virus scanning, child porn detection, etc.
- What other scanning do cloud providers do, and why? Do they tell you?
- Chris Hoofnagle (lawyer/law prof): “Google could use information it gleans from the messages for its own purposes—purposes it does not have to disclose to us”
- This is a matter of contract law—make sure you have a good, tech-knowledgeable lawyer read the contract

Other Cloud Applications

- There's no one answer to the security of other cloud applications—the answer is application-dependent
- The answer will depend on the security of the application, not of the concept of “cloud”
- In other words, we have to do a security assessment of each such application
- We can also use some heuristics: the inherent complexity of that application, the complexity of the configuration process, and the general reputation (if any) of the underlying software

- Web hosting *can* be very safe—web servers have had a lot of work done on hardening them—but there can still be configuration errors by the cloud customer, and there is still the problem of customer-written scripts
- But those two problems are no different in the Cloud!
- What about shared document editing?

A cloud VM takes as much system administration as a real computer running the same OS.

From Amazon's [Shared Responsibility Model](#):

“AWS operates, manages and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the service operates. The customer assumes responsibility and management of the guest operating system (including updates and security patches), other associated application software as well as the configuration of the AWS provided security group firewall.”

You save sysadmin effort if and only if you'd be running your own hypervisor host.

Secure Execution

- Microsoft's Azure Confidential Computing offers "Trusted Execution Environments"
- Using Intel SGX—but *properly* dividing a program into SGX/non-SGX pieces is *hard*
- (Still, that division is at the heart of iOS's excellent security architecture)
- Virtual Secure Mode: hypervisor-enforced protection against cloud administrators
- But—can a cloud administrator patch the hypervisor? What if Secure Boot is enabled? What if a Windows Hyper-V developer conspires with an Azure administrator?

Cloud-Specific Issues

- The provider
- The provider's employees
- Other customers
- The cloud platform

The Cloud Provider

- Do you trust the provider's honesty? Its competence?
- This is no different than any other business arrangement!
- How do you vet your suppliers of any other goods or services?
- N.B.: Many companies outsource janitorial and security services. These are sensitive, too
- But: be *certain* that you have a feasible way to migrate your data and applications to a new provider should that prove necessary

Google's Security Practices

- Internal vulnerability management
- Internal intrusion detection monitoring
- Incident response management
- Hardware tracking and disk erasure
- Regulatory compliance
- Multifactor authentication
- Internal traffic encryption

(From Google Cloud Security and Compliance Whitepaper)

The Provider's Employees

- Many of the same arguments apply to provider employees—but there's more that can be done here
- What are the technical controls that protect you from rogue employees?
Process controls?
- What sorts of behavior and log file auditing are done?
- Has any outside firm verified that these protections are actually in place?
That they work?
- Note well: what are your technical, procedural, etc., controls on your own system administrators?
- “Sed quis custodiet ipsos custodes?”

Google's Employee Security Measures

- Background checks
- Training
- Internal audit team
- Least privilege access; very limited physical access to hardware

Are other customers of the cloud provider a risk? Maybe, but it's hard

VM escape Hypervisors aren't bug-free. Can someone on the same VM break into the hypervisor and then use that access to penetrate your VM?

Meltdown, Spectre, etc. Some hardware attacks are threats if the enemy can run code on the same computer as you

Firewall break Different customers' VMs are firewalled from each other—can someone break the firewall?

- How do you configure your cloud service settings?
- How do you manage authentication and authorization?
- How do you manage sharing?
- Do you *thoroughly* understand the myriad options—probably via a web interface, rather than the code your engineers are accustomed to—and know how to set things up correctly?
- Remember the “open aws bucket” issue. . .

Amazon's Claim

With AWS, you control where your data is stored, who can access it, and what resources your organization is consuming at any given moment. Fine-grain identity and access controls combined with continuous monitoring for near real-time security information ensures that the right resources have the right access at all times, wherever your information is stored. Reduce risk as you scale by using our security automation and activity monitoring services to detect suspicious security events, like configuration changes, across your ecosystem. You can even integrate our services with your existing solutions to support existing workflows, streamline your operations, and simplify compliance reporting.

(From

<https://aws.amazon.com/security/>)

“You control” security—but do you know how to? “Fine-grained” gives you lots of control—but it’s a lot harder to get right. And do you know how to do the “integration” properly?

How Do We Analyze This?

- Cloud services have advantages and risks—and there is no one answer, because there is no one service
- How do we decide if we should use these services?
- As always: what are the benefits, and what are the risks to execution environments?

- Organizations *must* have email
- Almost certainly, a large cloud email provider can do it more securely than you can (and can do a better job of spam-filtering)
- The big risk: the contract terms, especially about privacy
- Note: news organizations that may want to write critical stories about these providers have a more complex calculus

- What are your *real* needs?
- Storage requirements that vary sharply over time?
- Storage shared *within* your organization?
- Storage available to telecommuters?
- Storage shared with other companies?
- (For the latter three, what are the relevant execution environments?)

Analyzing Cloud Storage

- If your storage needs vary sharply over time, that suggests that cloud storage is a good idea
- For internal use, you may not need it
- For telecommuters, would a VPN suffice? Do you have enough bandwidth? Low-enough latency compared with cloud storage providers?
- But—how do telecommuters synchronize files on laptops with file server copies? Dropbox does that well; can you?
- For sharing with external companies, good cloud services are probably easier—unless you're setting up a larger-scale shared complex
- Note carefully: we opt for less possible expansion of execution environments unless there's a large payoff

- Again: do your computing needs vary sharply over time?
- If not, what is the benefit of cloud computers?
- You probably have to administer the VMs' operating systems, just as in-house
- Easy virtual machine creation is an advantage—but you can do that yourself (though you may have to write your own management platform if non-privileged users should be able to spin up VMs)
- But—what do space and electricity cost? Do you save more, by moving your computers to a semi-remote area, than it would cost you to host things yourself, even with the incremental risk?
- But—how predictable is your budget? Can you count on funds to replace your server complex every few years, or are operational funds easier to get than capital funds?

Other Applications

- How important are these applications? How much would it cost you to maintain them?
- Example: payroll handling requires very specialized software, updated (at least) annually to account for tax law changes. You almost certainly do *not* write your own—and it's often easier to outsource all of it than to buy a new software package every year
- How often do you need real-time, shared document editing?
- If your documents tend to be solo or written by two-person teams, does simple screen-sharing suffice?
- Again: what is the (financial) benefit compared with the added risk?

- Generic concern: what is the availability of cloud services?
- Alternate question: what is the availability of your services? If a computer or router of yours fails, what is the fail-over? What is the repair time?
- Cloud services aren't perfect; their outages tend to make the news
- But—what is the average, aggregated customer downtime for cloud user, compared with the aggregated downtime of roll-your-own, in-house services?

Bottom Line

- For large companies, most services can be done in-house easily enough, *if* they devote enough resources and attention to system administration
- For smaller companies, cloud services make a lot of sense—outsource as much as you can and concentrate on your own business. Most of the big cloud companies are honest and competent—but make sure you avoid provider lock-in
- Also: you don't have to decide on all or nothing; you can use cloud services for some needs while keeping other stuff in-house

Questions?



(Great blue heron, Central Park, March 23, 2021)