CS 4530: Fundamentals of Software Engineering Lesson 5.1 Cloud Infrastructure

Rob Simmons Khoury College of Computer Sciences

© 2025 Released under the <u>CC BY-SA</u> license

Learning objectives for this lesson

- By the end of this lesson, you should be able to...
 - Explain what "cloud" computing is and why it is important
 - Explain why shared infrastructure is important in cloud computing
 - Discuss trade-offs that you might consider for self or vendor-managed platforms

Old school



Data centers were the original shared infrastructure

- Save on physical costs (power, cooling, security)
- More cost-effective reliable network connections

Old school



• Comcast might get mad?

Old school

closet in Jamacia Plain



Also handles HTTPS certificates

Old school problems



My web server that I bought



The raspberry pi in my utility closet in Jamacia Plain

- Servers were (are?) expensive
- Being in a different place than your server is annoying
- Very hard to respond to changes in demand

Infrastructure as a Service (IaaS)

- Let the people running the data center handle the physical computers (purchasing, maintaining, replacing)
- Required new software and systems, especially virtualization which lets you run multiple independent operating systems on one computer (mid-2000s)
- Linode and DigitalOcean in 2003, Amazon's Elastic Compute Cloud (EC2) launches in 2006
- Fast (minutes/hours not days/weeks) to get more capacity

Platform as a Service (PaaS)

Amazon Web Services launched with three products:

- EC2 virtual machine hosting is *infrastructure* as a service.
- Simple Queue Service (SQS) handled the problem of needing to do too many things at once (oversimplification: avoid running out of memory)
- Simple Storage Service (S3) handled the problem of needing to store and deliver "static" data (oversimplification: avoid running out of disk space and/or network capacity)
- SQS and S3 are *"platforms"* or basic capabilities
- Typically very fast (minutes or fractions of a second) to get more capacity

Example: my Mastodon instance



X as a Service

- Infrastructure as a Service bring-your-own-operatingsystem, but don't worry about the hardware
- Platform as a Service hello, here's a microservice or database or file storage location for you to use
- Software as a Service here's Microsoft Word on the internet for you to use

MongoDB Atlas. Fully managed MongoDB in the cloud.

Gray areas exist — are MongoDB Atlas and other DBaaS services PaaS or Saas? — but useful distinctions in practice.

Cloud Infrastructure is best for variable workloads

- Consider:
 - Does your workload benefit from ability to scale up or down?
 - Variable workloads have different demands over time (most common)
 - Constant workloads require sustained resources (less common)
- Example:
 - Need to run 300 VMs, each 4 vCPUs, 16GB RAM
- Private cloud:
 - Dell PowerEdge Pricing (AMD EPYC 64 core CPUs)
 - 7 servers, each 128 cores, 512GB RAM, 3 TB storage = \$162,104
- Public cloud:
 - Amazon EC2 Pricing (M7a.xlarge instances, \$0.153/VM-hour)
 - 10 VMs for 1 year + 290 VMs for 1 month: \$45,792.90
 - 300 VMs for 1 year: \$402,084.00

Public clouds are not the only option

- "Public" clouds are connected to the internet and available for anyone to use
 - Examples: Amazon, Azure, Google Cloud, DigitalOcean
- "Private" clouds use cloud technologies with onpremises, self-managed hardware
 - Cost-effective when a large scale of baseline resources are needed
 - Example management software: OpenStack, VMWare, Proxmox, Kubernetes
- "Hybrid" clouds integrate private and public (or multiple public) clouds
 - Effective approach to "burst" capacity from private cloud to public cloud
- Remember to think about your threat model





The CLOUD Act and Transatlantic Trust



https://www.csis.org/analysis/cloud-act-and-transatlantic-trust

"X as a Service" offers several abstractions to choose from depending on your needs

- Vendor manages different levels of the stack, achieving economies of scale
- When would you choose one over the other?
- Explore some options at <u>https://compareclou</u> <u>d.in/</u>



Learning objectives for this lesson

- By the end of this lesson, you should be able to...
 - Explain what "cloud" computing is and why it is important
 - Explain why shared infrastructure is important in cloud computing
 - Discuss trade-offs that you might consider for self or vendor-managed platforms