

Virtualization/Vagrant/ Cloud Computing

Amol Deshpande

University of Maryland

amol@cs.umd.edu

Cloud Computing

- ▶ Computing as a “service” rather than a “product”
 - Everything happens in the “cloud”: both storage and computing
 - Personal devices (laptops/tablets) simply interact with the cloud
- ▶ Advantages
 - Device agnostic – can seamlessly move from one device to other
 - Efficiency/scalability: programming frameworks allow easy scalability (relatively speaking)
 - Increasing need to handle “Big Data”
 - Reliability
 - Cost: “pay as you go” allows renting computing resources as needed – much cheaper than building your own systems

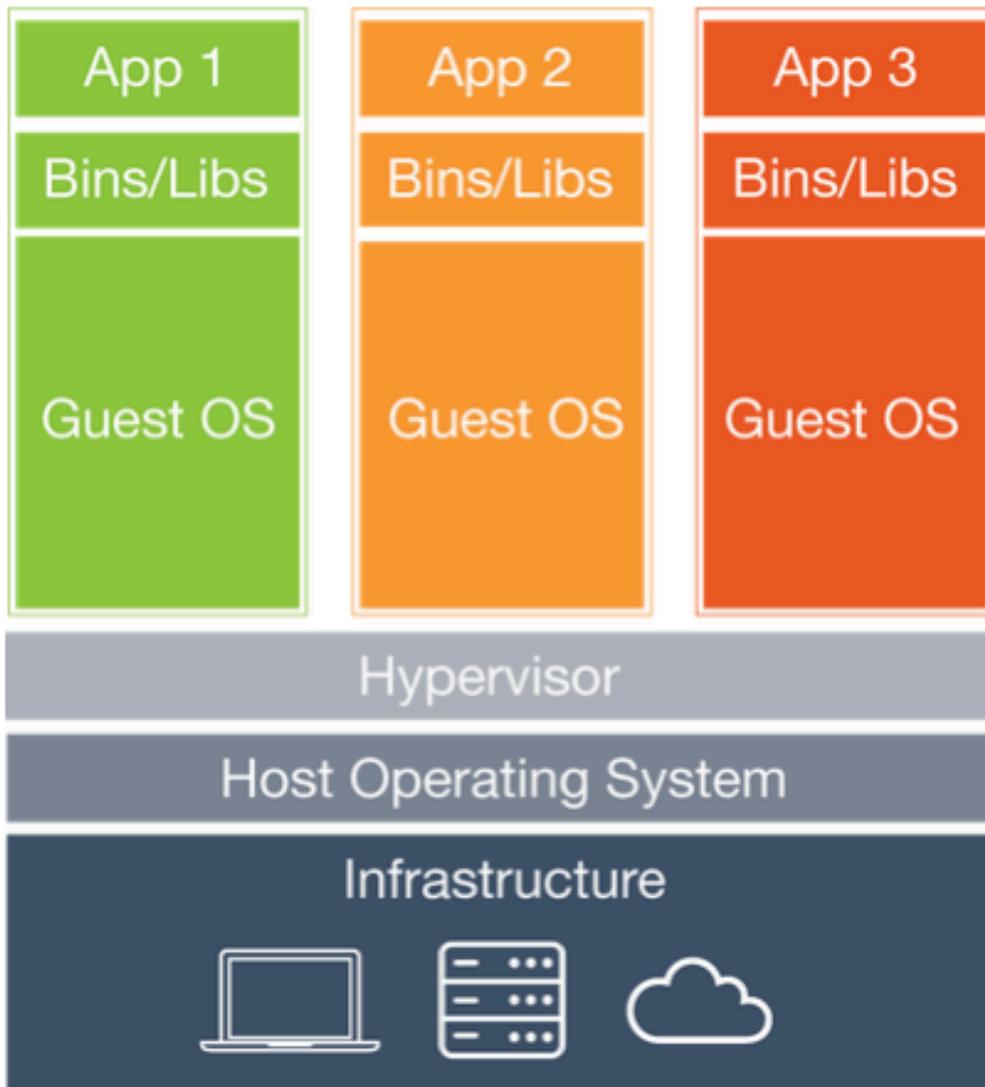
Cloud Computing

- ▶ Basic ideas have been around for a long time (going back to 1960's)
 - Mainframes + thin clients (more by necessity)
 - Grid computing a few year ago
 - Peer-to-peer
 - Client-server models
 - ...
- ▶ But it finally works as we wished for...
 - Why now?... A convergence of several key pieces over the last few years
 - Does it really? ... Still many growing pains

Virtualization

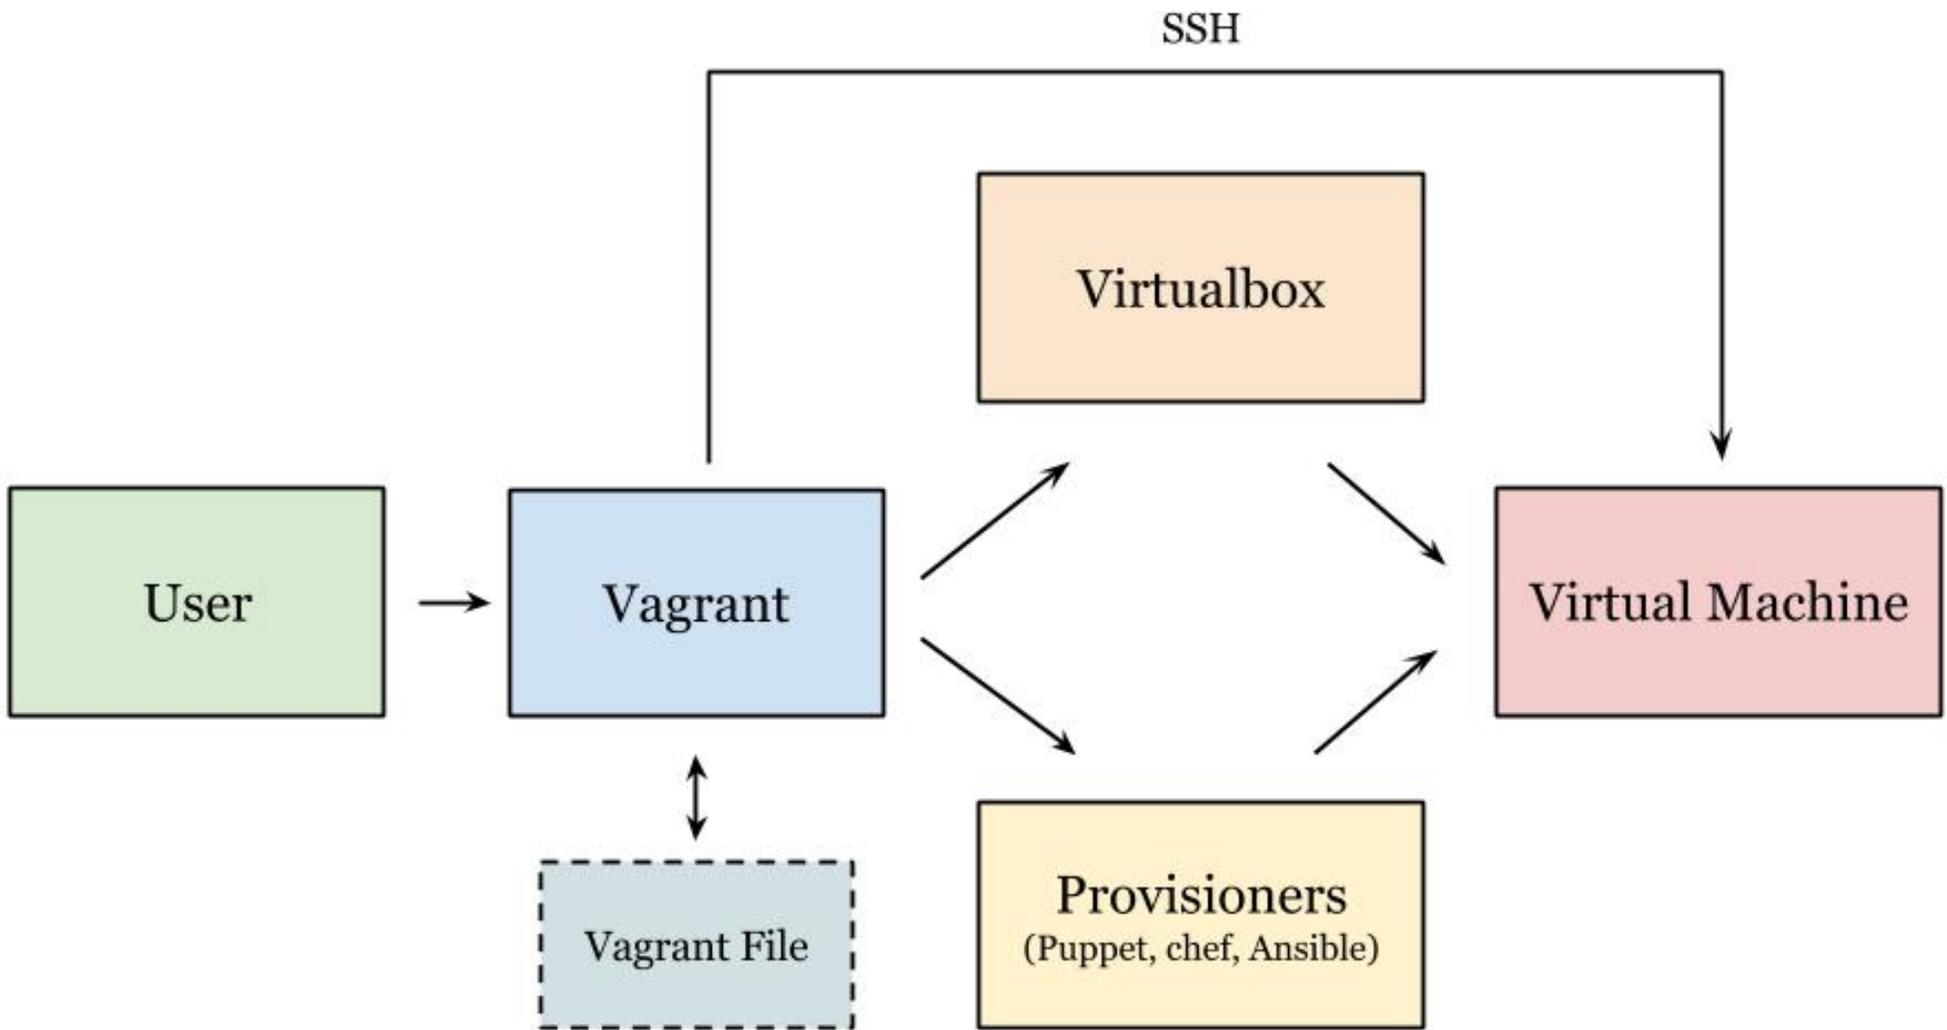
- ▶ Virtual machines (e.g., running Windows inside a Mac) etc. has been around for a long time
 - Used to be very slow...
 - Only recently became efficient enough to make it a key for CC
- ▶ Basic idea: run virtual machines on your servers and sell time on them
 - That's how Amazon EC2 runs
- ▶ Many advantages:
 - Security: virtual machines serves as almost impenetrable boundary
 - Multi-tenancy: can have multiple VMs on the same server
 - Efficiency: replace many underpowered machines with a few high-power machines

Virtual Machines

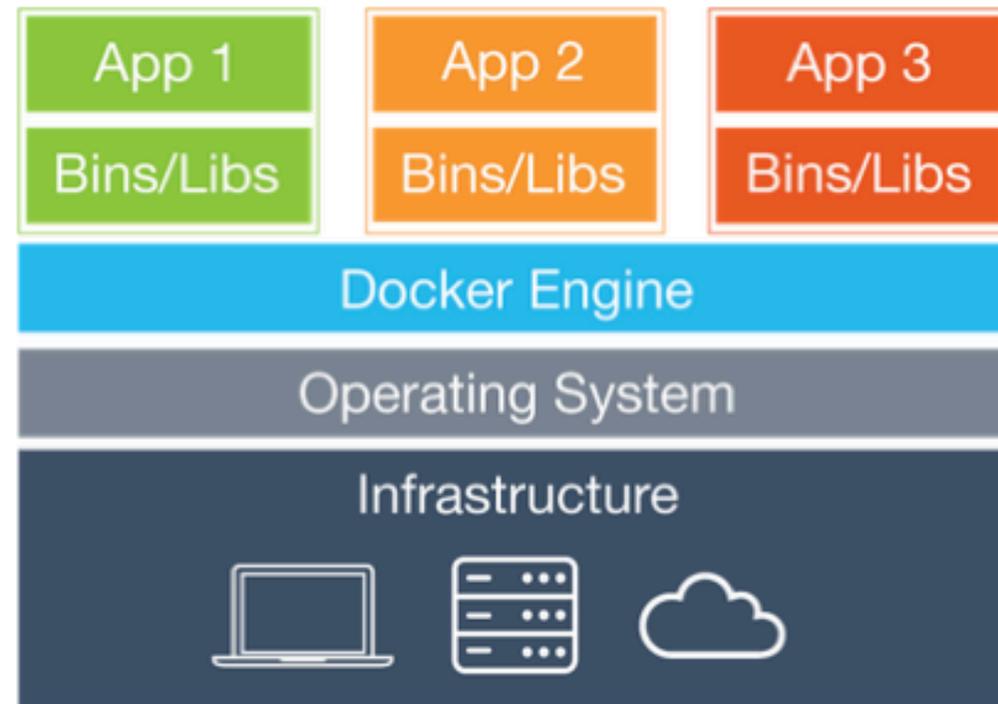
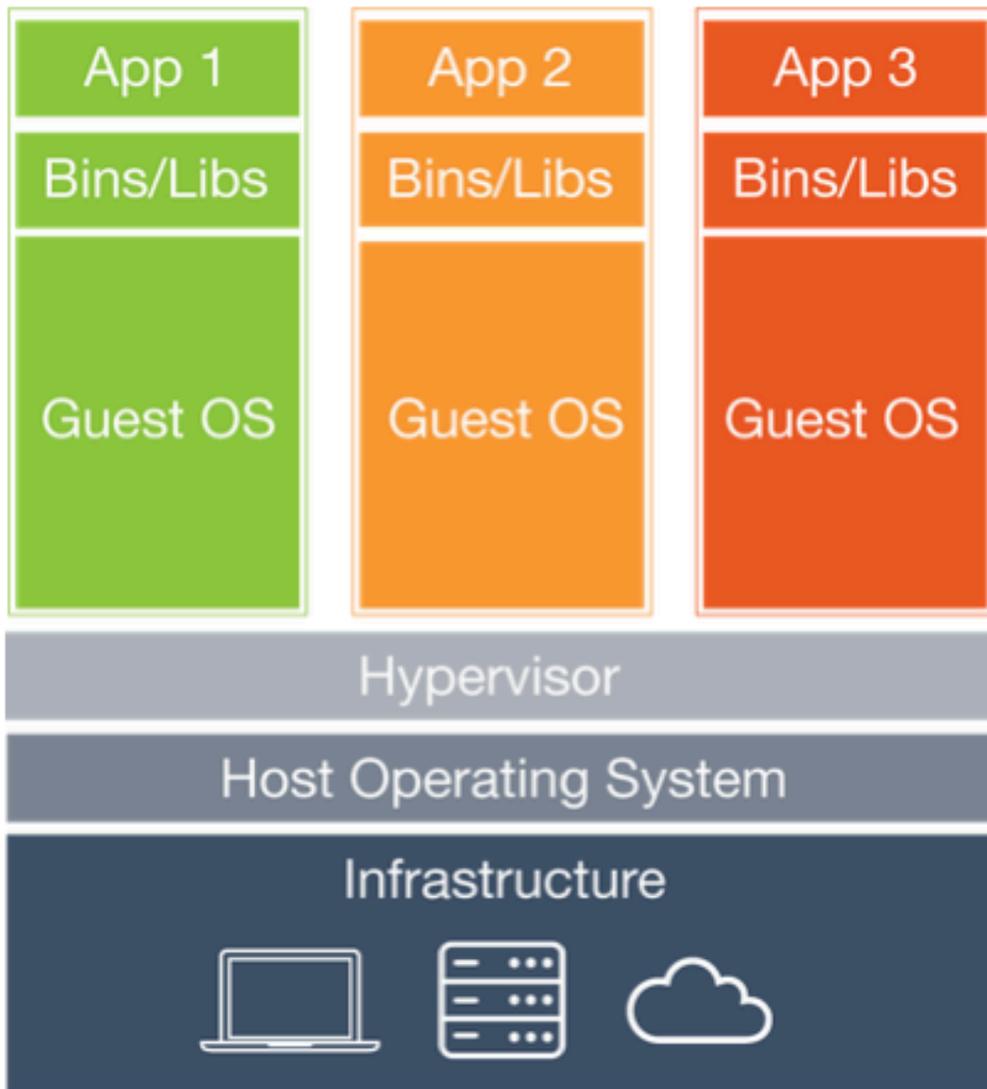


Vagrant

Vagrant makes it easy to create and configure virtual environments.



Virtual Machines vs Containers



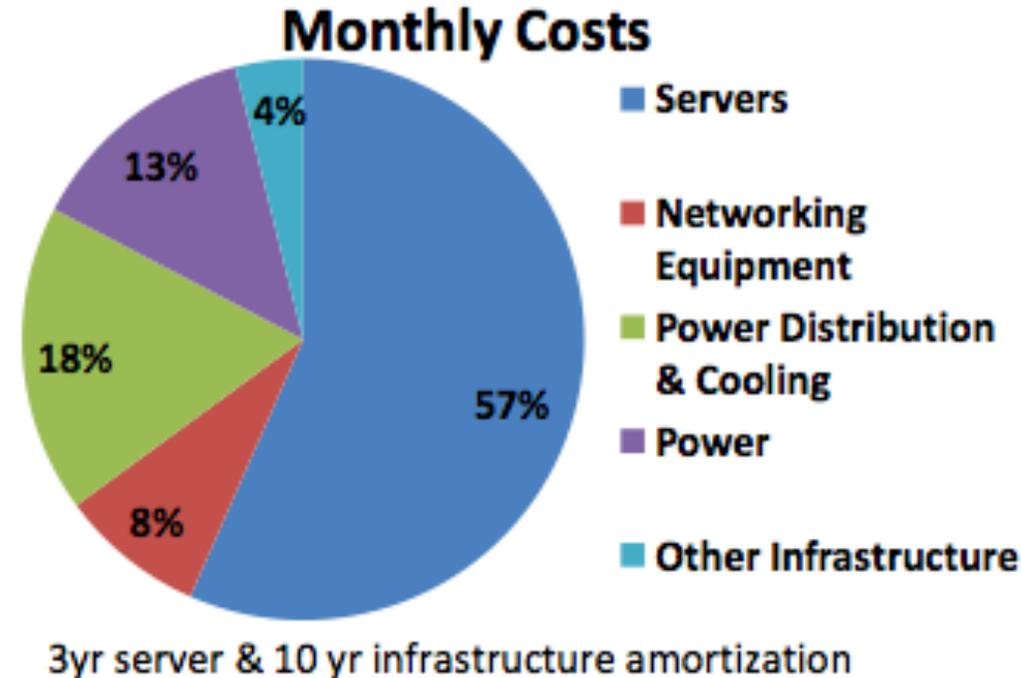
Data Centers

- ▶ The key infrastructure piece that enables CC
- ▶ Everyone is building them
- ▶ Huge amount of work on deciding how to build/design them



Data Centers

- ▶ Amazon data centers:
Some recent data
 - 8 MW data center can include about 46,000 servers
 - Costs about \$88 million to build (just the facility)
 - Power a pretty large portion, but server costs still dominate



“Every day, Amazon Web Services adds enough new capacity to support all of Amazon.com’s global infrastructure through the company’s first 5 years, when it was a \$2.76B annual revenue enterprise”

Putting it together

- ▶ A Cloud Computing Provider builds and manages “Data Centers”, often with millions of servers
- ▶ They may rent you:
 - Raw hardware (not that common)
 - Virtual machines in those data centers (Infrastructure-as-a-service)
 - A software service that does something specific for you (Software-as-a-service)
 - Something in-between (Platform-as-a-service)



On-Premises



IaaS

Infrastructure as a Service



PaaS

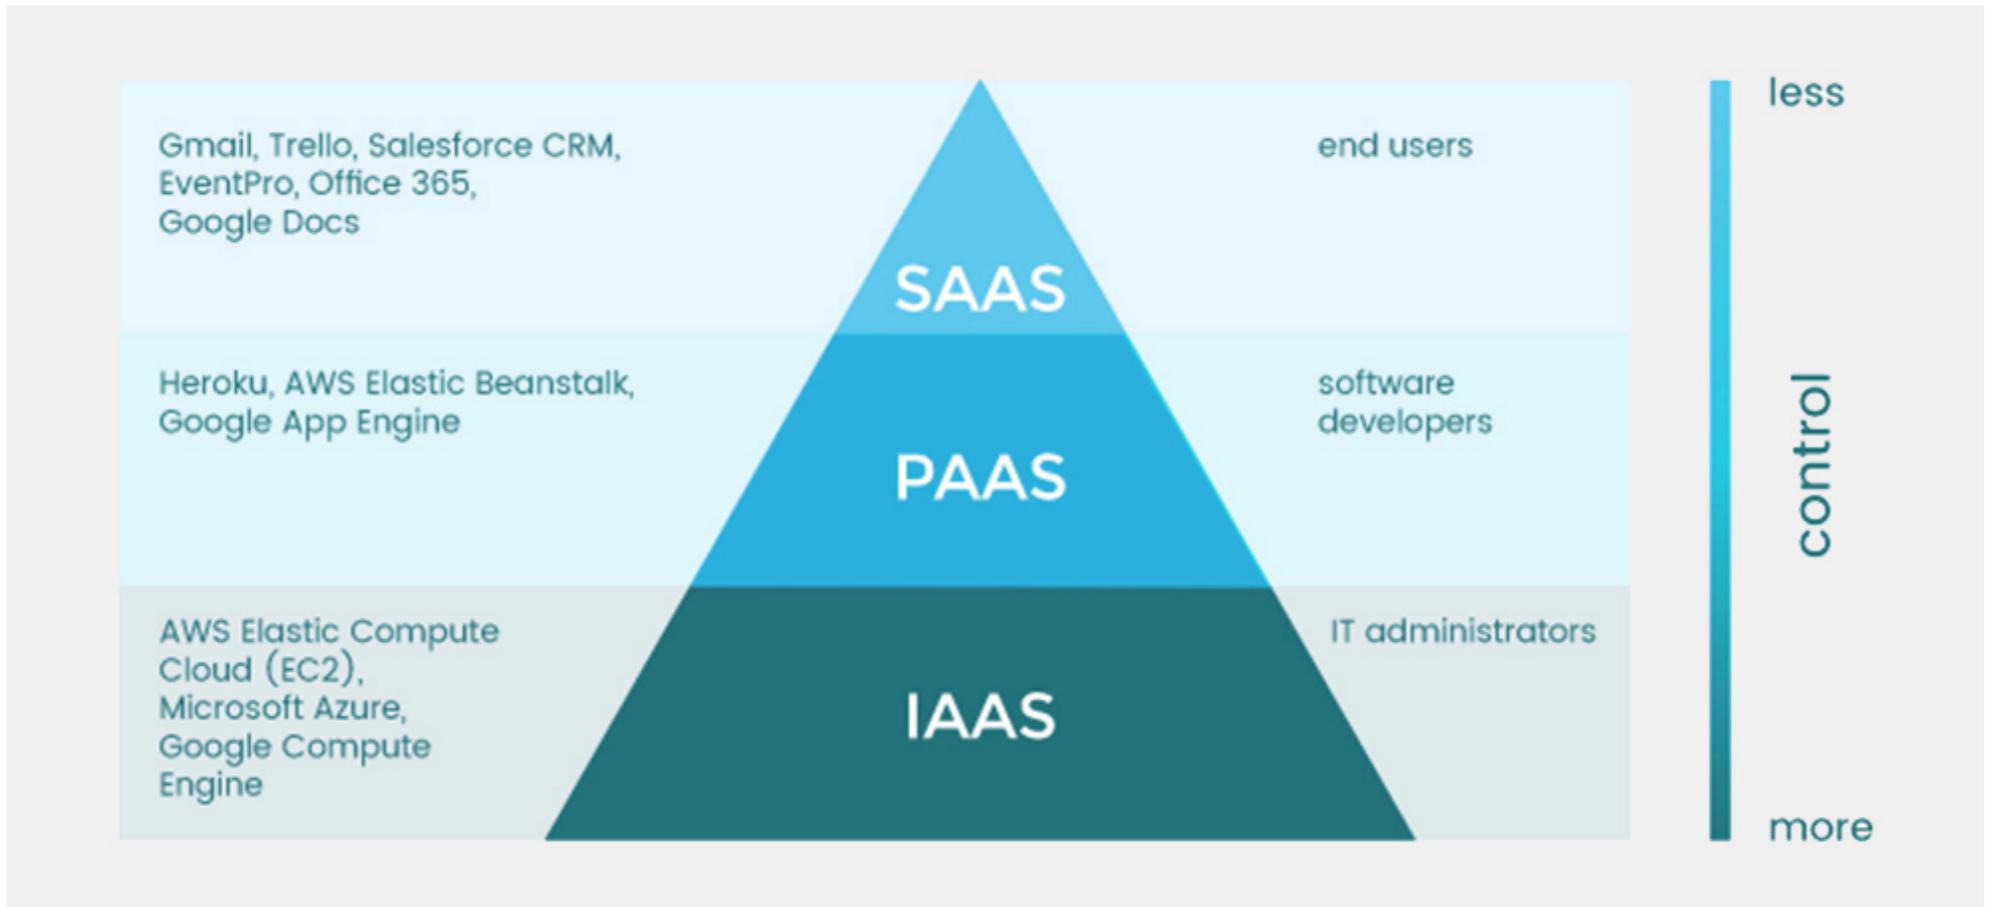
Platform as a Service



SaaS

Software as a Service

On-Premises	IaaS	PaaS	SaaS
Applications	Applications	Applications	Applications
Data	Data	Data	Data
Runtime	Runtime	Runtime	Runtime
Middleware	Middleware	Middleware	Middleware
O/S	O/S	O/S	O/S
Virtualization	Virtualization	Virtualization	Virtualization
Servers	Servers	Servers	Servers
Storage	Storage	Storage	Storage
Networking	Networking	Networking	Networking

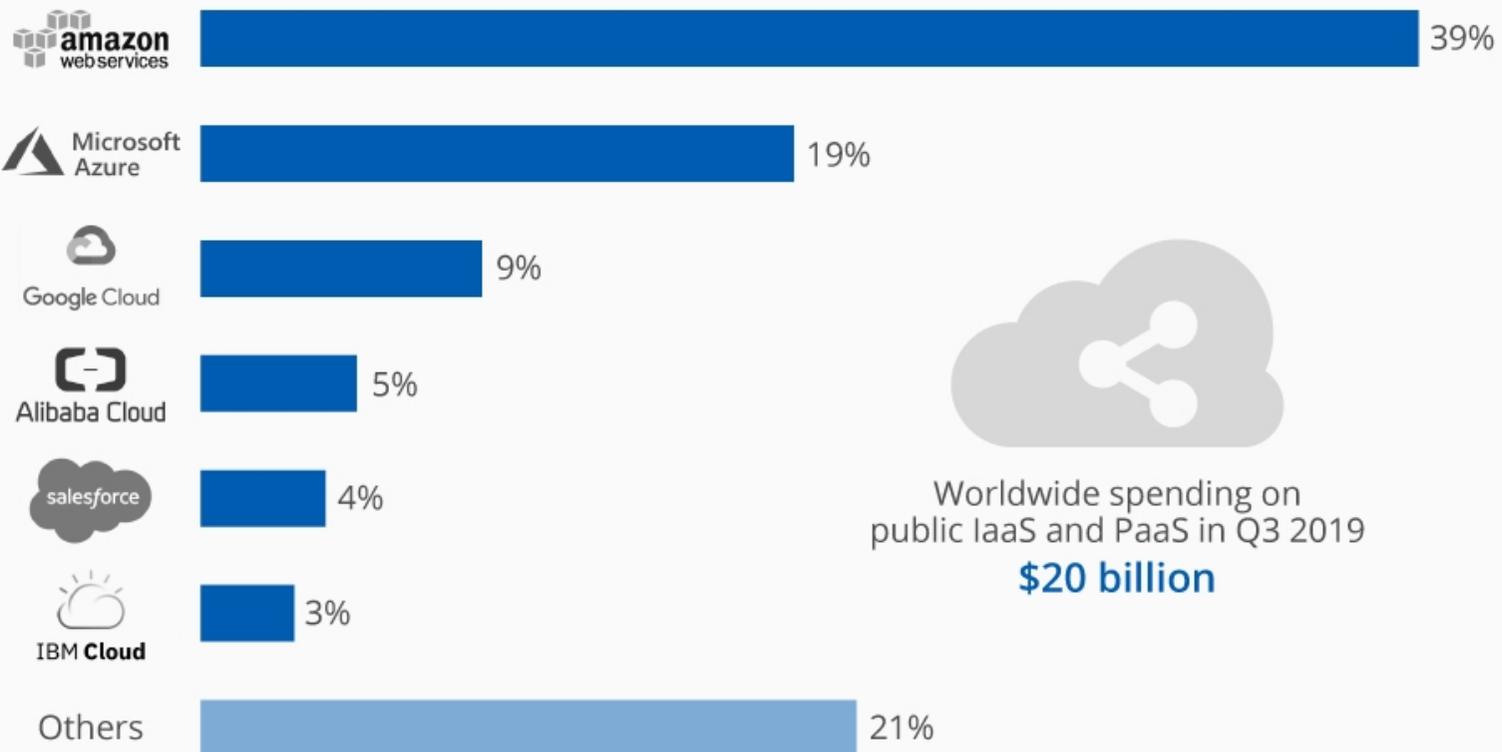


Amazon Web Services

- ▶ Leader in this space for a while

Amazon Dominates Public Cloud Market

Global market share of public cloud infrastructure service providers in Q3 2019*



AWS has 175 Services Today.. More every year...

Compute

- Amazon EC2
- AWS Lambda
- Amazon EC2 Container Service
- Elastic Load Balancing
- AWS Elastic Beanstalk
- Auto Scaling

Network

- Amazon VPC
- Amazon Route 53
- AWS Direct Connect

Storage

- Amazon S3
- Amazon CloudFront
- Amazon Glacier
- Amazon Elastic File System
- AWS Storage Gateway
- AWS Import/Export Snowball

Security & Identity

- AWS Identity and Access Management
- AWS KMS
- AWS Directory Service
- AWS Cloud HSM
- AWS WAF

Applications

- Amazon WorkDocs
- Amazon WorkSpaces
- Amazon WorkMail

Databases

- Amazon RDS
- Amazon DynamoDB
- Amazon ElastiCache
- Amazon Redshift
- AWS Database Migration Service

Analytics

- Amazon EMR
- AWS Data Pipeline
- Amazon Elasticsearch Service
- Amazon Machine Learning
- Amazon Kinesis

App Services

- Amazon SES
- Amazon AppStream
- Amazon SWF
- Amazon Elastic Transcoder
- Amazon CloudSearch
- Amazon SQS
- Amazon API Gateway

Management Tools

- AWS CloudFormation
- AWS Config
- AWS CloudTrail
- AWS Service Catalog
- AWS OpsWorks
- Amazon CloudWatch
- Trusted Advisor
- AWS Certificate Manager

Developer Tools

- AWS CodeCommit
- AWS CodeDeploy
- AWS CodePipeline

Mobile Services

- Amazon Cognito
- AWS Device Farm
- Amazon SNS
- Amazon Mobile Analytics
- Mobile Hub

Internet of Things

- AWS IoT

Microsoft Azure not far behind

Platform Services

Security & Management

- Security Center
- Portal
- Azure Active Directory
- Azure AD B2C
- Multi-Factor Authentication
- Automation
- Scheduler
- Key Vault
- Store/Marketplace
- VM Image Gallery & VM Depot

Media & CDN

- Media Services
- Media Analytics
- Content Delivery Network

Integration

- API Management
- BizTalk Services
- Logic Apps
- Service Bus

Compute Services

- Container Service
- VM Scale Sets
- Batch
- RemoteApp
- Dev/Test Lab

Application Platform

- Web Apps
- Mobile Apps
- API Apps
- Cloud Services
- Service Fabric
- Notification Hubs
- Functions

Developer Services

- Visual Studio
- Mobile Engagement
- VS Team Services
- Xamarin
- Application Insights
- HockeyApp

Data

- SQL Database
- SQL Data Warehouse
- DocumentDB
- SQL Server Stretch Database
- Redis Cache
- Storage Tables
- Azure Search

Intelligence

- Cognitive Services
- Bot Framework
- Cortana

Analytics & IoT

- HDInsight
- Machine Learning
- Stream Analytics
- Data Catalog
- Data Lake Analytics Service
- Data Lake Store
- IoT Hub
- Event Hubs
- Data Factory
- Power BI Embedded

Hybrid Cloud

- Azure AD Health Monitoring
- AD Privileged Identity Management
- Domain Services
- Backup
- Operational Analytics
- Import/Export
- Azure Site Recovery
- StorSimple

Infrastructure Services

Compute

- Virtual Machines
- Containers

Storage

- Blob
- Queues
- Files
- Disks

Networking

- Virtual Network
- Load Balancer
- DNS
- Express Route
- Traffic Manager
- VPN Gateway
- App Gateway

Datacenter Infrastructure

