VM Migration and Cloud Computing Overview

Yiying Zhang

Outline

- Project proposal graded. AWS awaiting university.

- VM migration: “*Live Migration of Virtual Machines*”

- What is cloud computing?
- Why cloud computing?
- How clouds are built?

The Need for VM Migration

- Load balancing
  - Move VMs from heavily loaded to less loaded host machines
  - Move VMs to new machines added to the cluster
- Recovery from host failure
  - Move (restart) VMs on a new host when its original host fails
- Maintenance of host
  - Temporarily move VMs off a host when this host is under maintenance

Types of Migration

- Whole VM migration
- Process/container migration

- Cold migration: shutdown a VM and restart it on a new host
- Warm migration: suspend a VM, move its states, resumes on a new host
- Live migration: move a VM’s states while it continues to run (with a brief stop)

Live VM Migration

- Migrate a running VM (without or with minimal disturbance to its applications)
  - All its in-memory, network, storage states
- Benefits
  - After migrating the whole VM, nothing left at the old server
  - In-memory states can be transferred consistently (and efficiently)
  - Allows a separation of concerns b/w the users and operator

Goals of Live VM Migration

- Minimize *downtime*
  - The time when services provided by VM are unavailable
- Small total *migration time*
  - The total time taken for all VM states to be migrated
- Migration has minimal effect on foreground services
  - Minimal resource contention with foreground VM services

Naive Migration: Checkpoint, Move, and Resume

- Pause the VM and take a snapshot of it
  - Record CPU registers, memory content, I/O states, file system
- Move the snapshot to the new machine
- Restore the snapshot
  - Restore in-memory states
- Resume the VM execution

Naive Migration Performance

- Time to move memory to memory + time to move disk/SSD to disk/SSD
- e.g., 10Gbps network, 120 MB/s disk sequential read/write bandwidth (7200 RPM disk), 17 GB/s (DDR4)
  - Migrating a VM with 16 GB memory and 32 GB disk
  - Time to move disk data: 32 GB / (120MB/s) = 4.55 min
  - Time to move memory data: 16 GB / 10Gbps = 12.8 sec
  - Plus the time to stop and resume the VM execution

➡ Downtime at least 4.6 min

Improving Storage Migration

- Use disaggregated storage (e.g., EC2, EBS)
- No need to move storage when migrating VMs
- What about local storage?
  - Only copy deltas of a base image?
  - Don’t copy everything all together?
  - Replicate local storage on multiple nodes (for reliability but could also benefit migration)

Improving Memory Migration

- Even if we don’t need to move storage, moving memory with the naive approach would still cause too much downtime (e.g., 12.8s for 16GB in our example)
- The focus of this paper
- This paper’s approach: ***pre-copy* migration**

Memory Migration Phases

- Push phase
  - Source VM continues running while pushing memory to new VM
- Stop-and-copy phase
  - Stop the source VM, copy all (remaining) memory to new VM, start new VM
- Pull phase
  - After new VM starts, pull memory from source VM as needed (accessed)

- Tradeoffs between these options?

Pre-Copy Migration

- Iteratively pre-copy (modified memory)
  - All memory copied in the first round
  - *n* round copies the modified memory content of round *n-1*
- A final step of stop-and-copy follows the pre-copy phase
  - This step determines downtime

- When shall we stop the pre-copy phase and move to stop-and-copy?

Writable Working Set

- For memory that is never/rarely modified (or at least not/rarely modified during the migration time) => pre-copy
- For memory that is frequently modified, i.e., *Writable Working Set* (*WWS*) => stop-and-copy
- Goal: estimate WWS to minimize the memory need to go through stop-and-copy (without making the pre-copy phase too long (affects total migration time))
- Could use heuristics or more smart ways to determine WWS

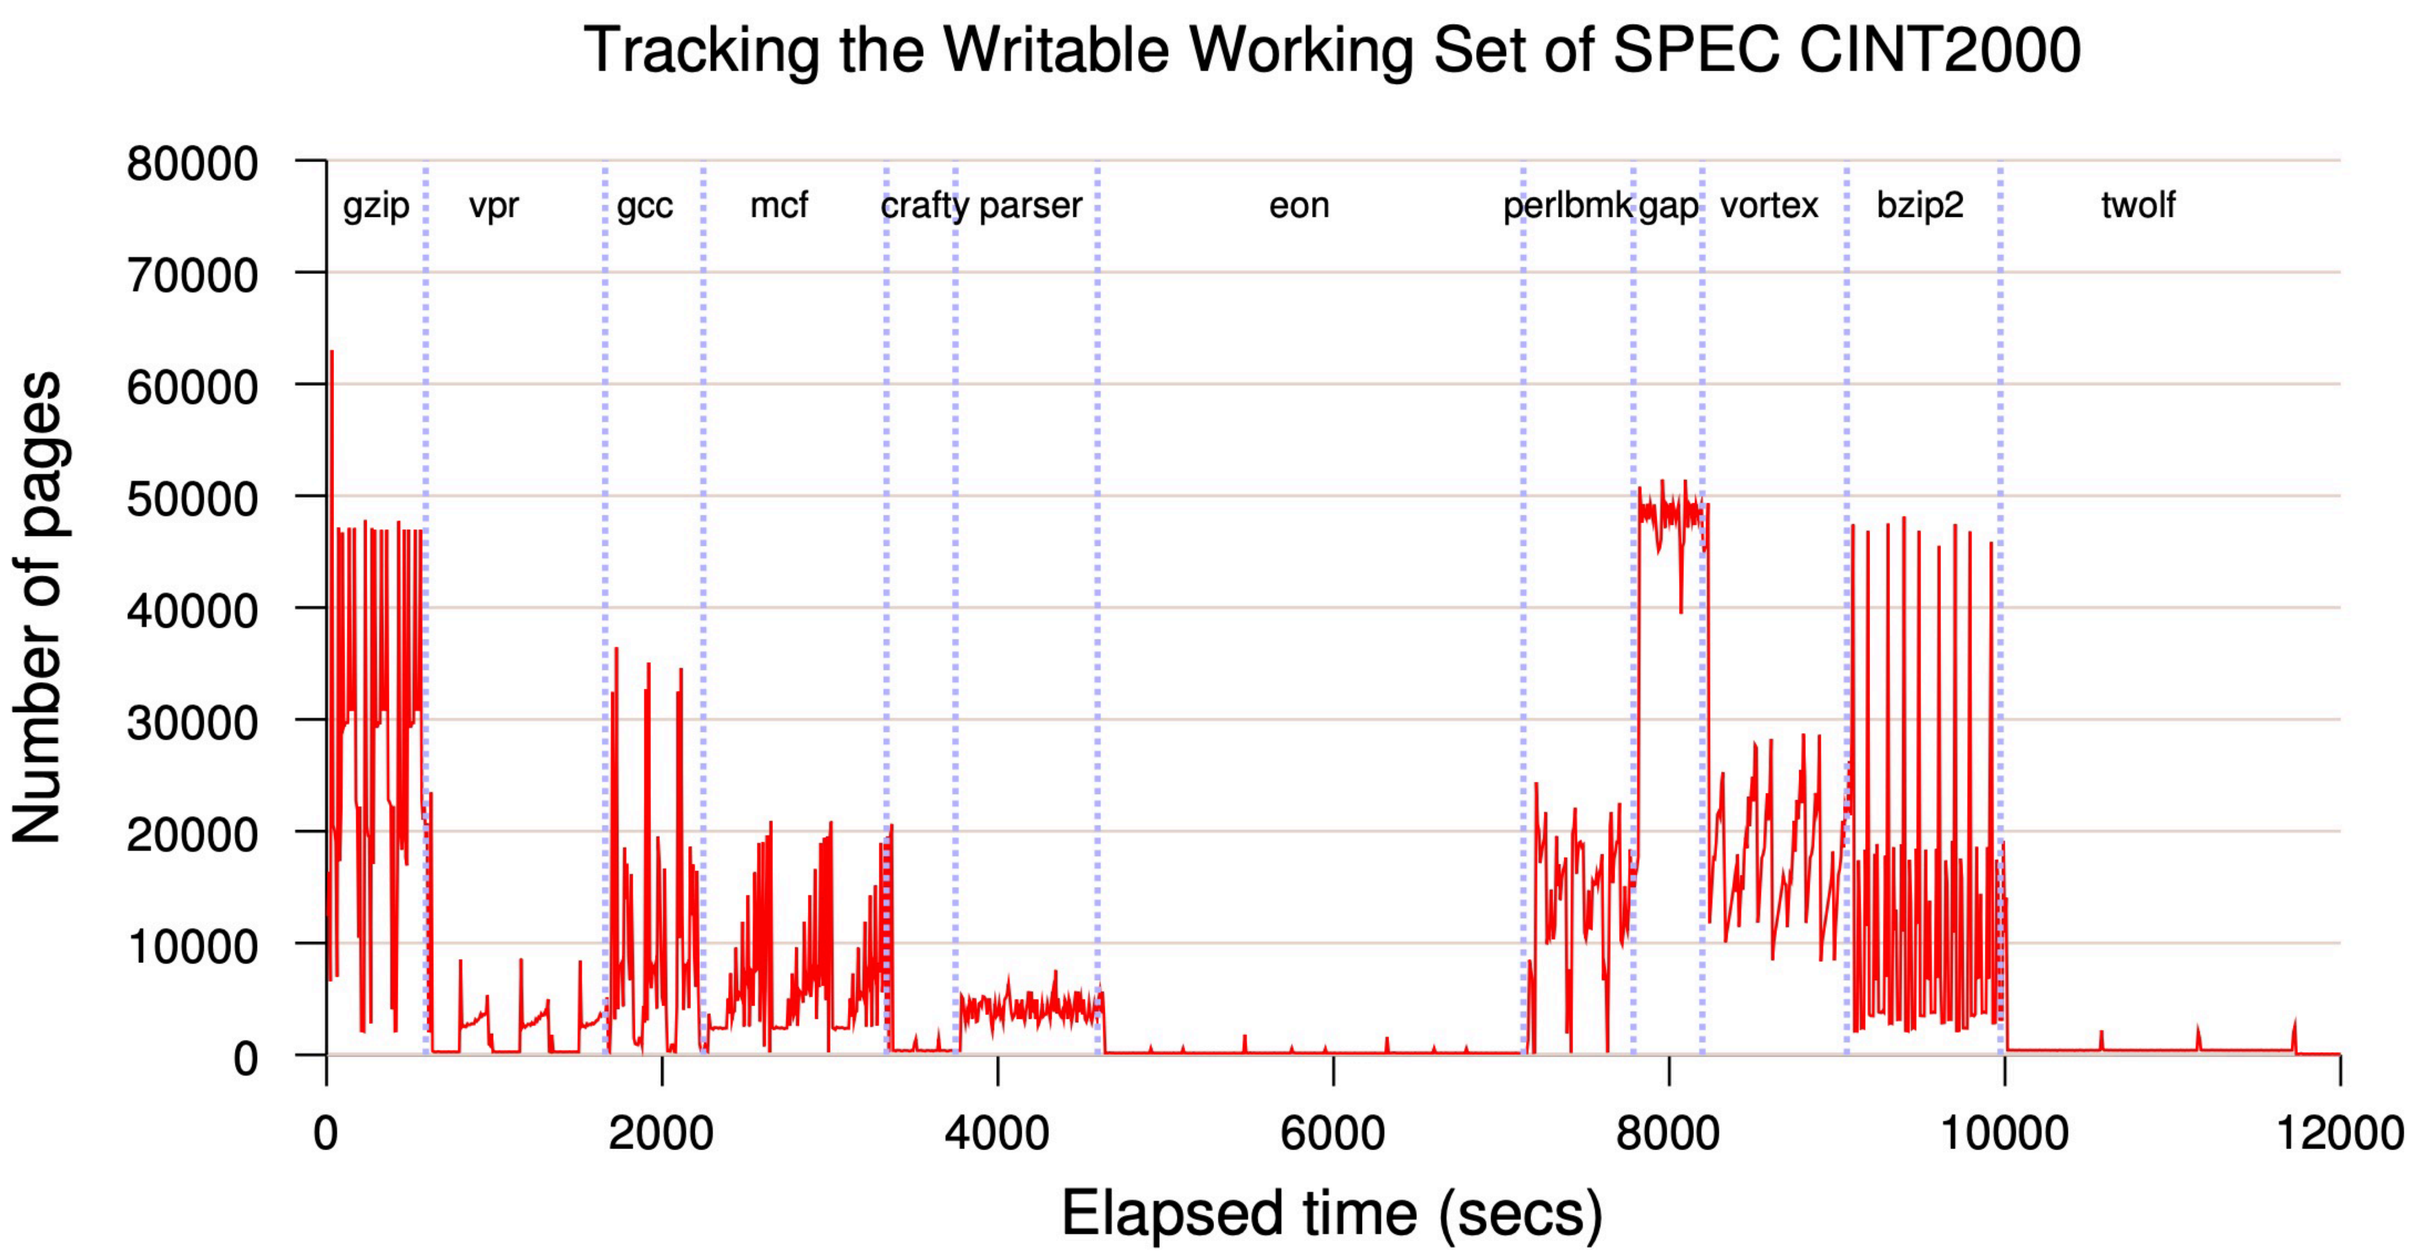


Figure 2: WWS curve for a complete run of SPEC CINT2000 (512MB VM)

Migrating Network State

- Goal: maintain all open network connections after migration without redirection
- Environment: migration within the same LAN
- Solution:
  - VM carries the original IP address and all protocol state
  - Generate unsolicited ARP (Address Resolution Protocol) reply from the migrated host to other nodes, advertising the new physical address of the IP
  - Other nodes reconfigure to send to the new physical address

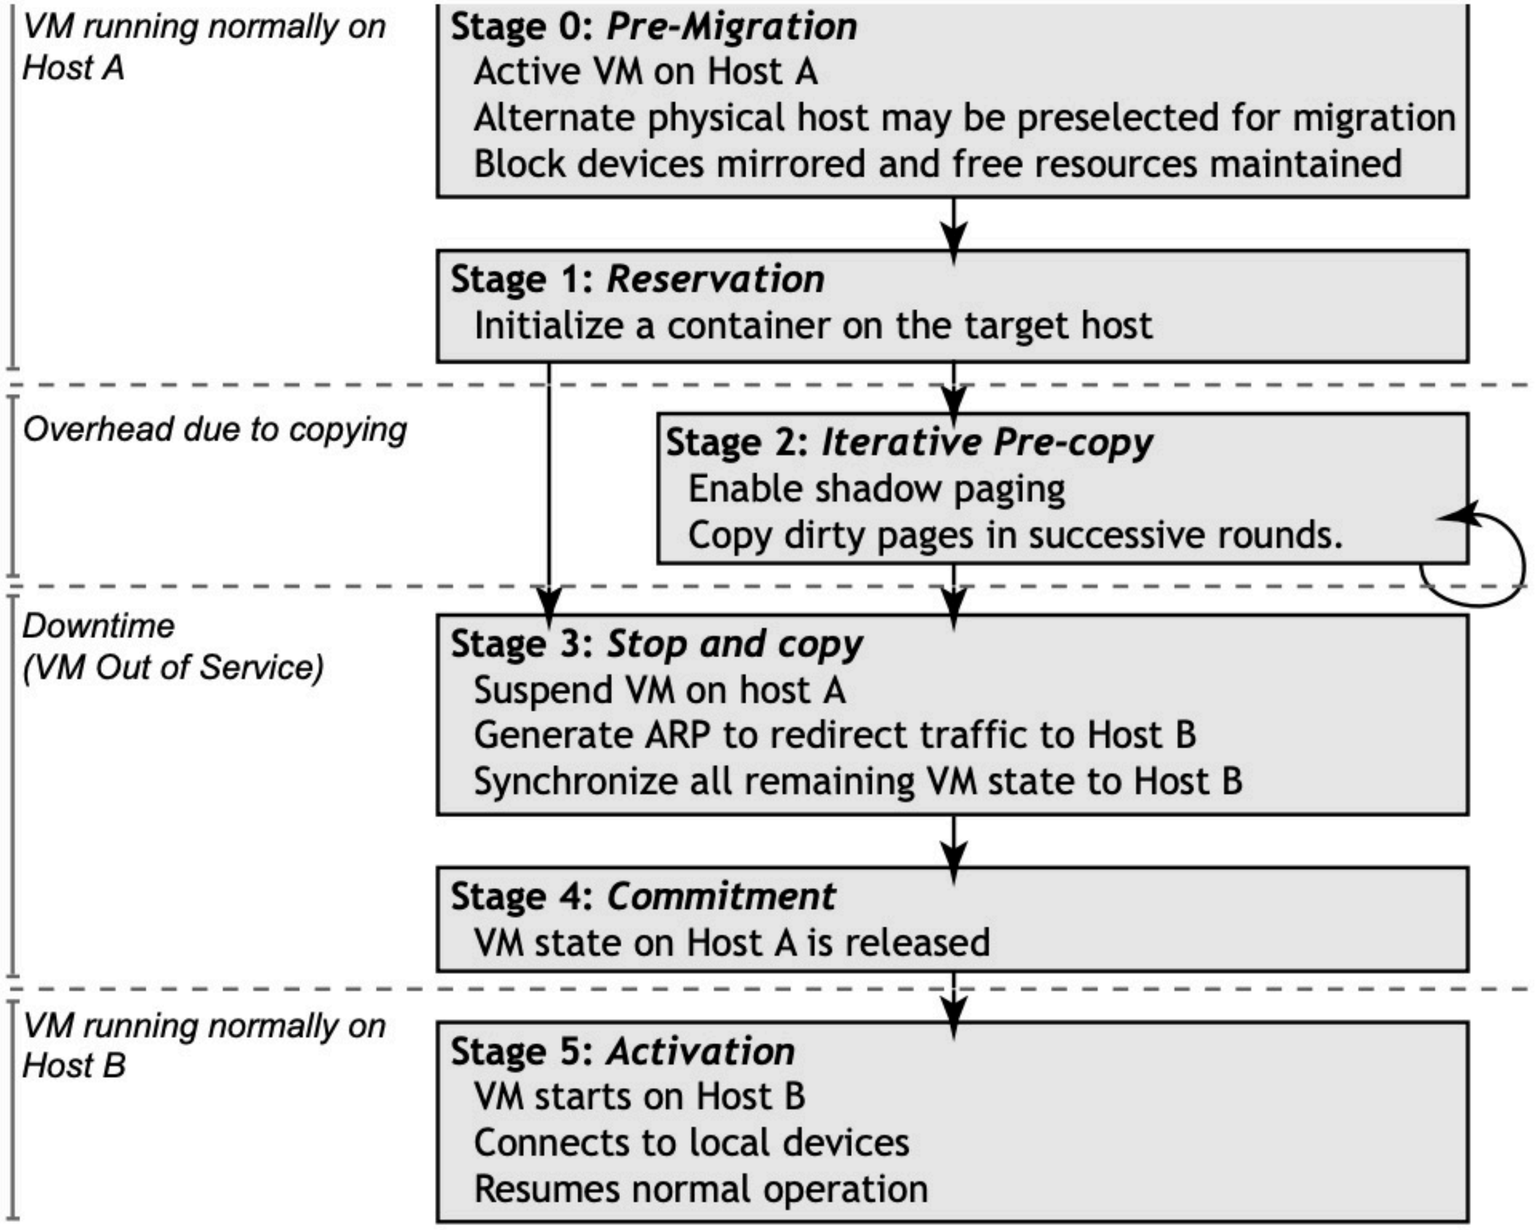


Figure 1: Migration timeline

Migration Mechanisms

- Managed migration
  - Migration performed by management VM (e.g., Dom0), guest OS agnostic
  - Needs to track dirty pages during rounds of pre-copy
  - Using an initially read-only shadow page tables in Xen to capture writes
- Self migration
  - Guest OS handles most of the migration
  - Guest OS tracks dirty memory and pre-copies them
  - Uses a two-stage stop-and-copy phase

Other Implementation Details

- Dynamic rate-limiting
  - Dynamically adapt the network bandwidth to balance between good migration performance and low impact on foreground activities
- Rapid page dirtying
  - Avoid transferring pages that are repeatedly dirtied
- Paravirtualized optimizations
  - Run a thread in VM to monitor the WWS of guest processes
  - Let guest OS free less useful memory (e.g., cold buffer cache pages)

Results

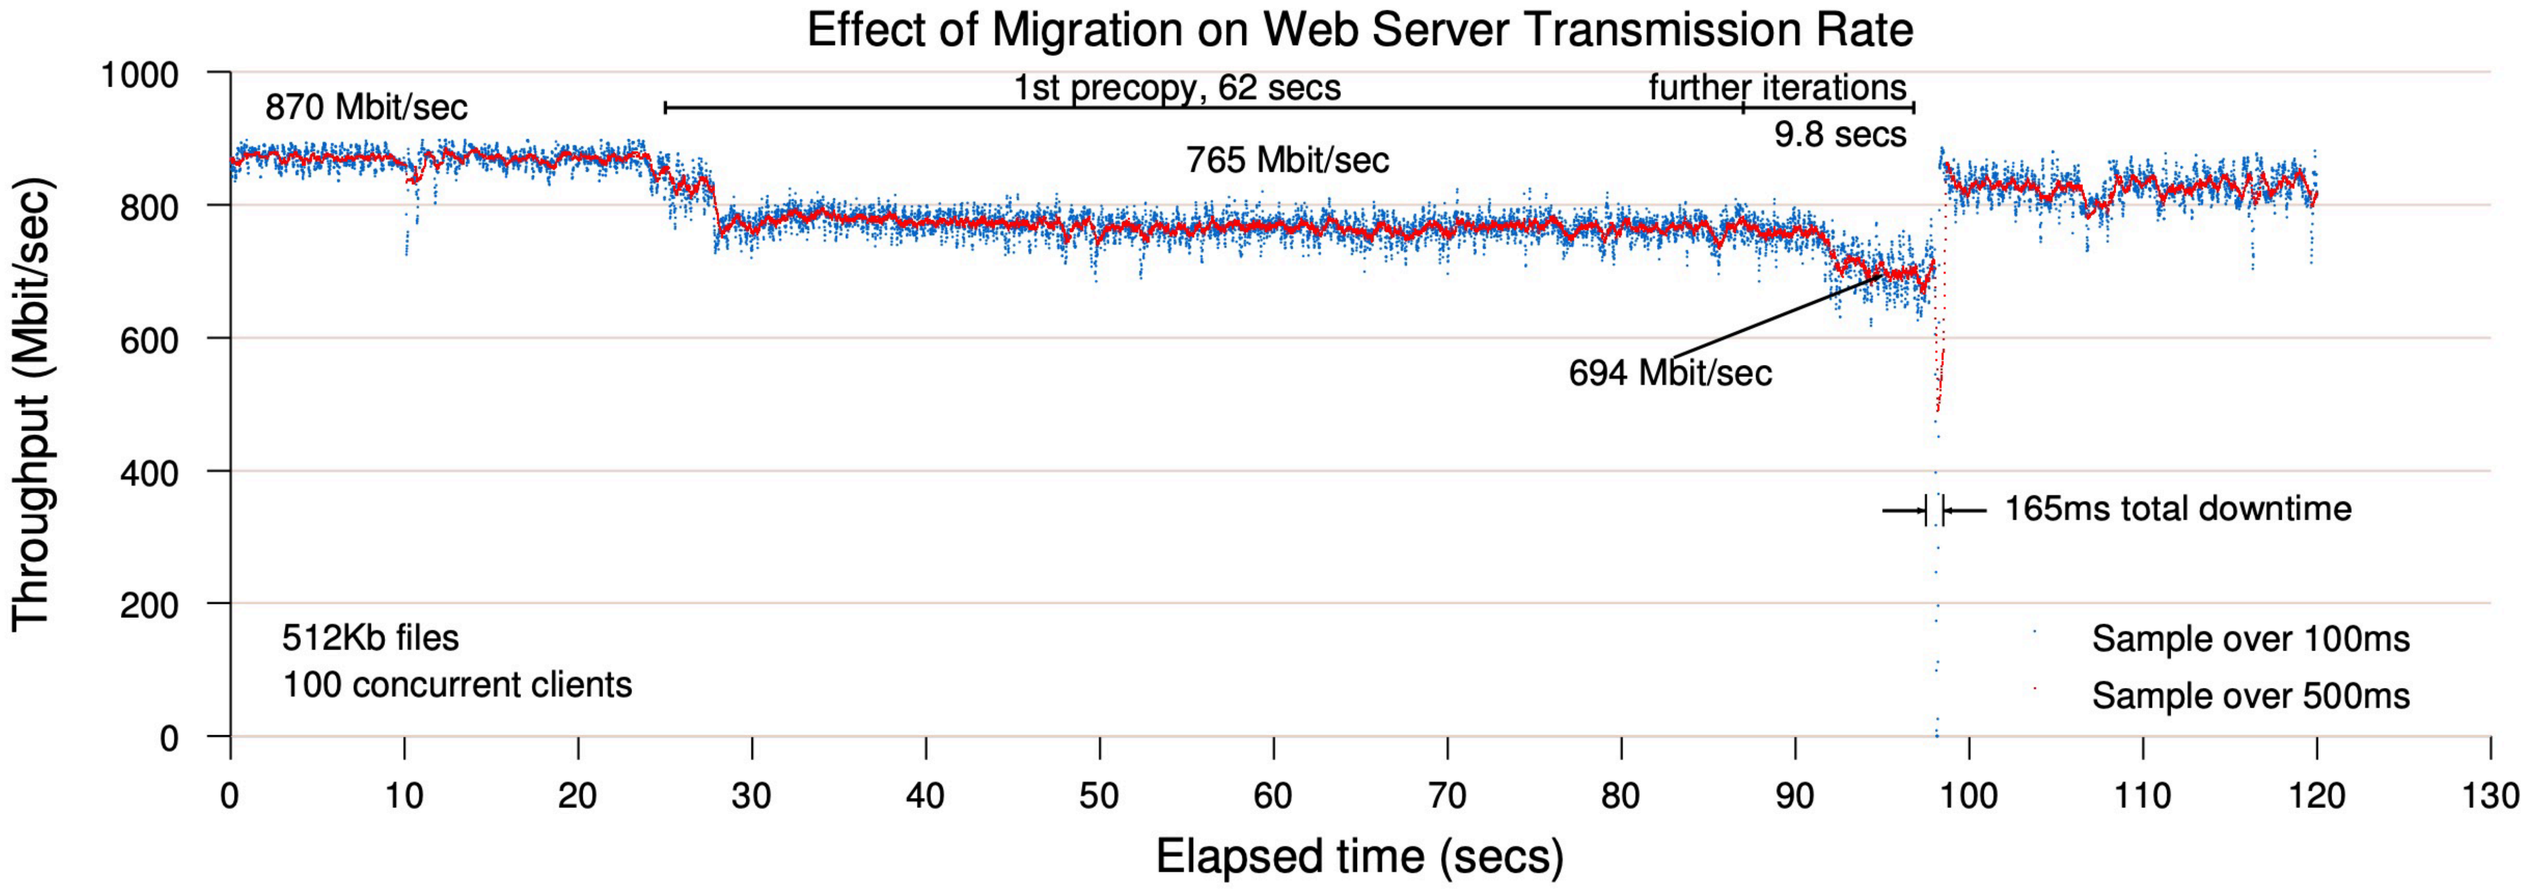


Figure 8: Results of migrating a running web server VM.

Results

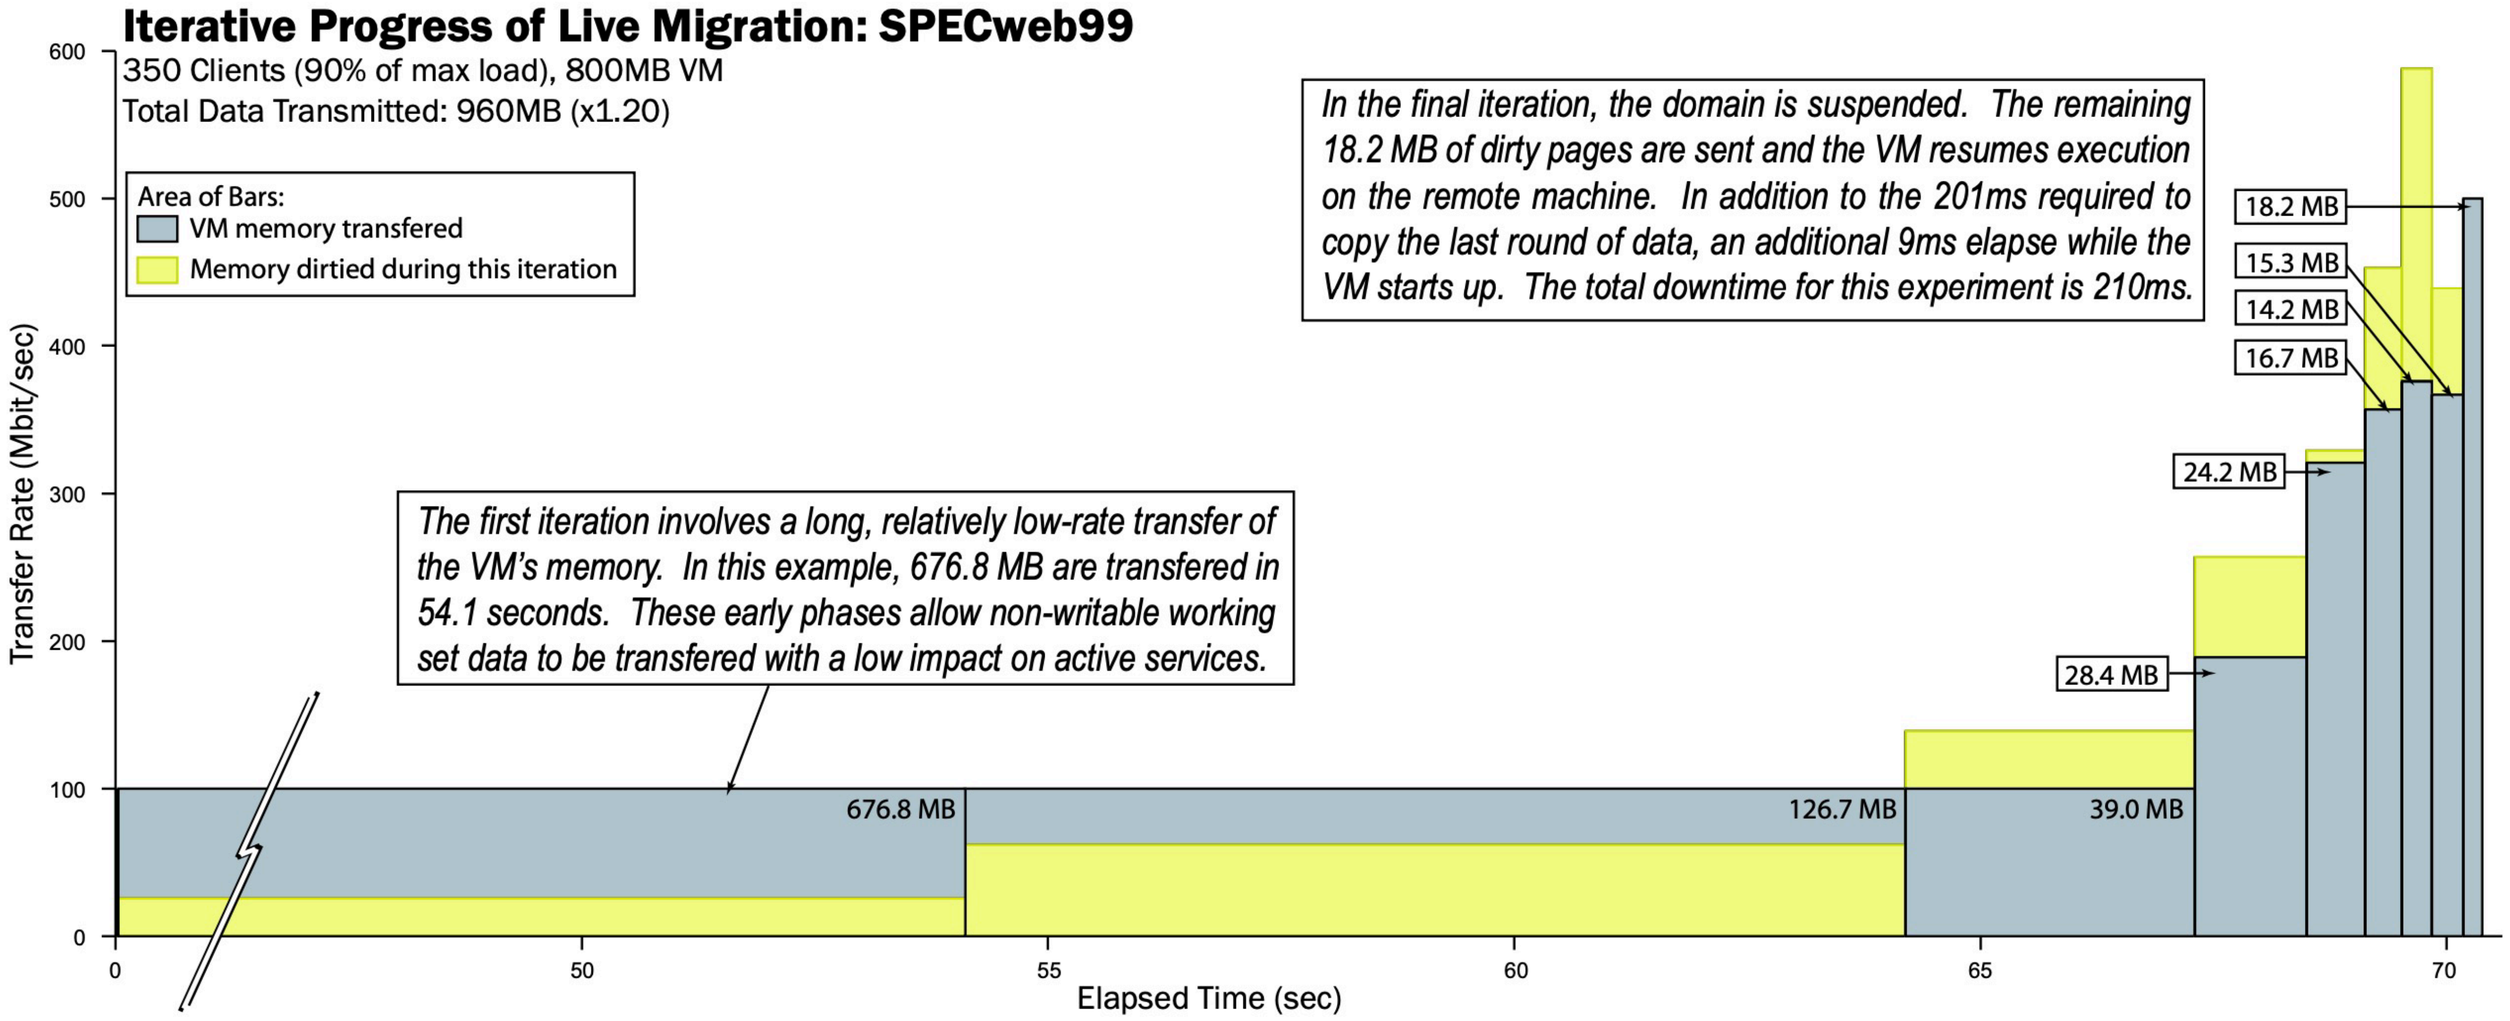


Figure 9: Results of migrating a running SPECweb VM.

Other VM Migration Systems

- VMware vMotion
- XenMotion
- Migration across heterogeneous architecture
- Migration across datacenters
- Migration from one cloud provider to another provider

Outline

- What is cloud computing?
- Why cloud computing?
- How clouds are built?

Cloud Computing

- Datacenters that rent servers or other computing resources (e.g., storage)
  - Anyone (or company) with a “credit card” can rent
  - Cloud resources owned and operated by a third-party (cloud provider)
- Fine-grain pricing model
  - Rent resources by time or by I/O
  - Pay as you go (pay for only what you use)
- Can vary capacity as needed
  - No need to build your own IT infrastructure for peak needs

Cloud Computing

1. The illusion of infinite computing resources available on demand

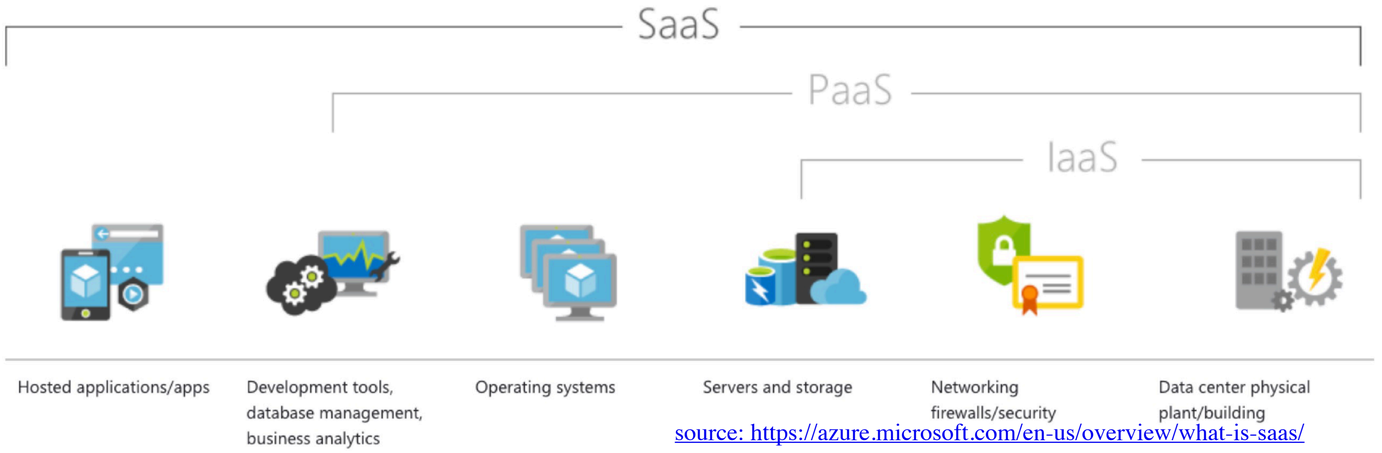
2. The elimination of an up-front commitment by cloud users

3. The ability to pay for use of computing resources on a short-term basis as needed

Source: Above the Clouds: A Berkeley View of Cloud Computing

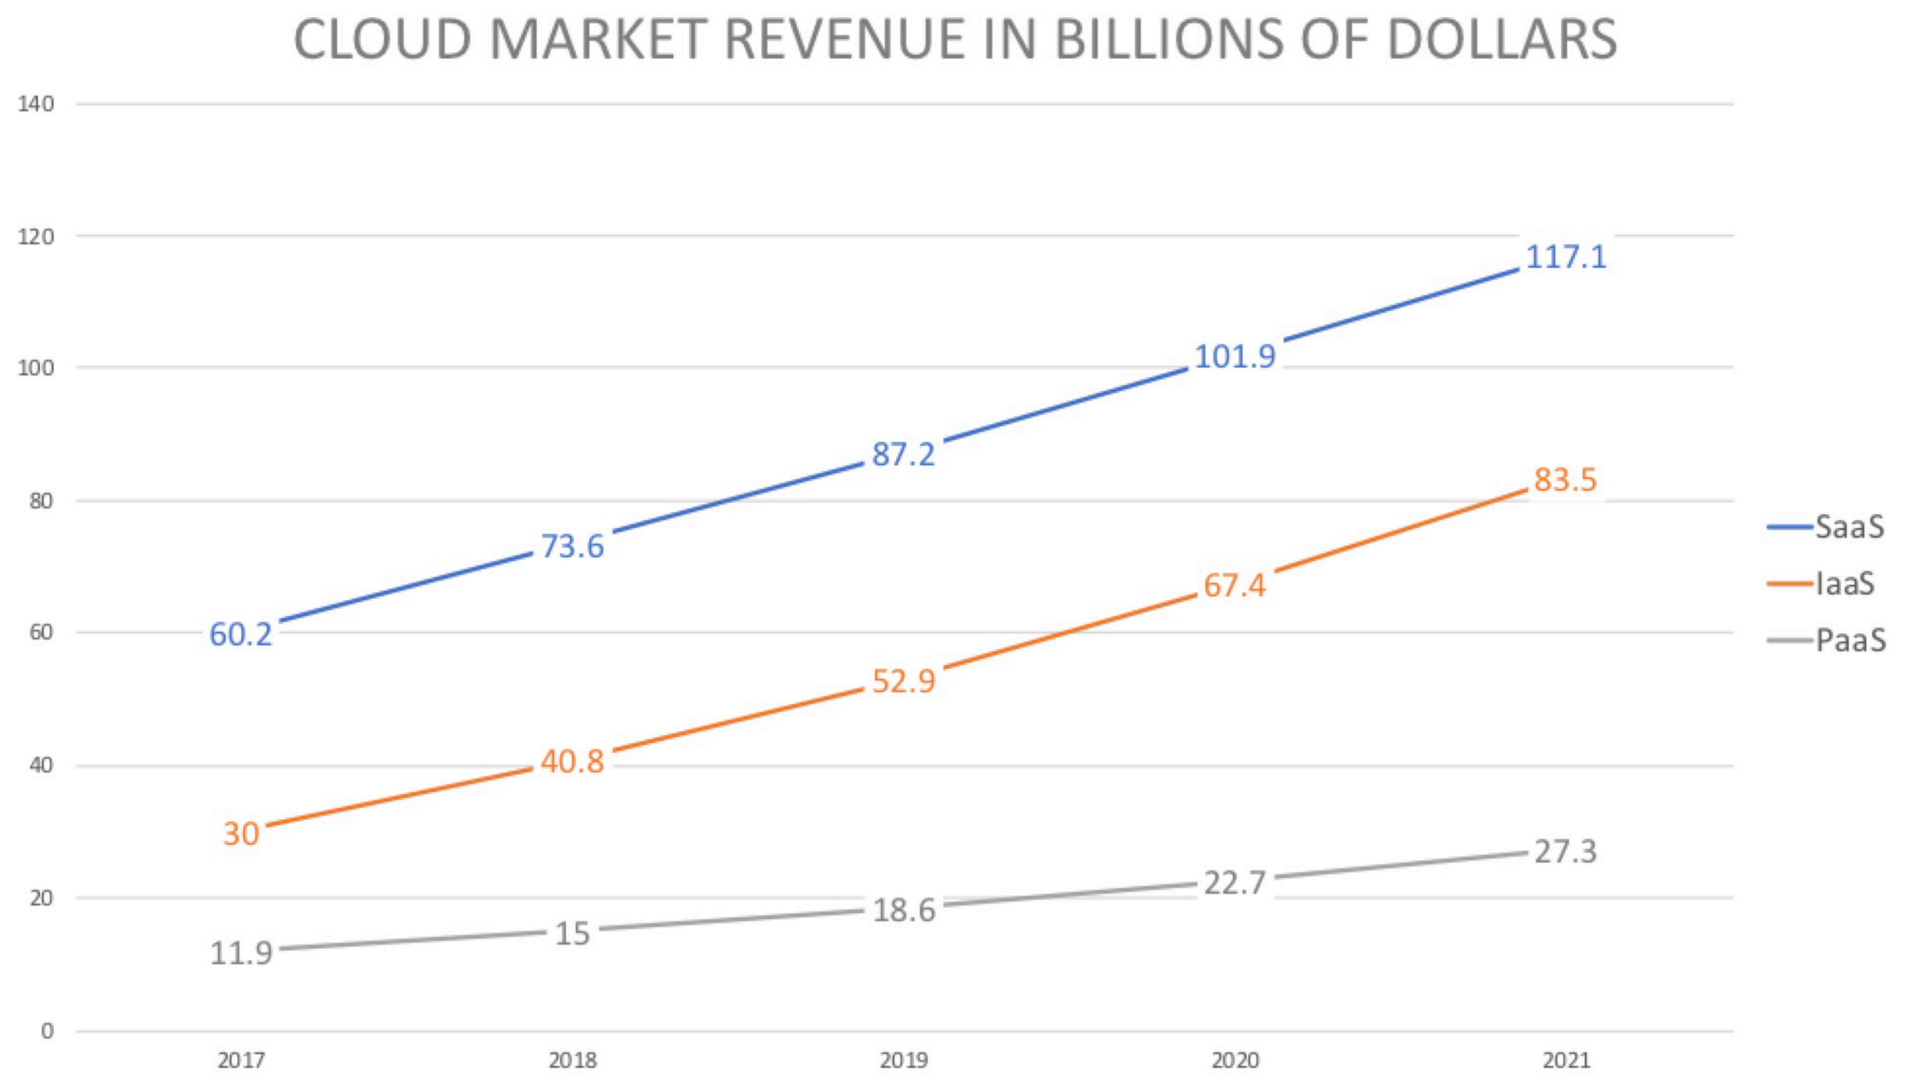
XaaS (what can be rented?)

- IaaS: Infrastructure as a Service
  - Sell VMs or physical servers
- PaaS: Platform as a Service
  - Platforms (environments) for developers to develop and run applications. e.g., Google App Engine
- SaaS: Software as a Service
  - Offer services/applications e.g., Salesforce, Snowflake
- FaaS: Function as a Service
- All can be deployed at (public) cloud or local datacenters



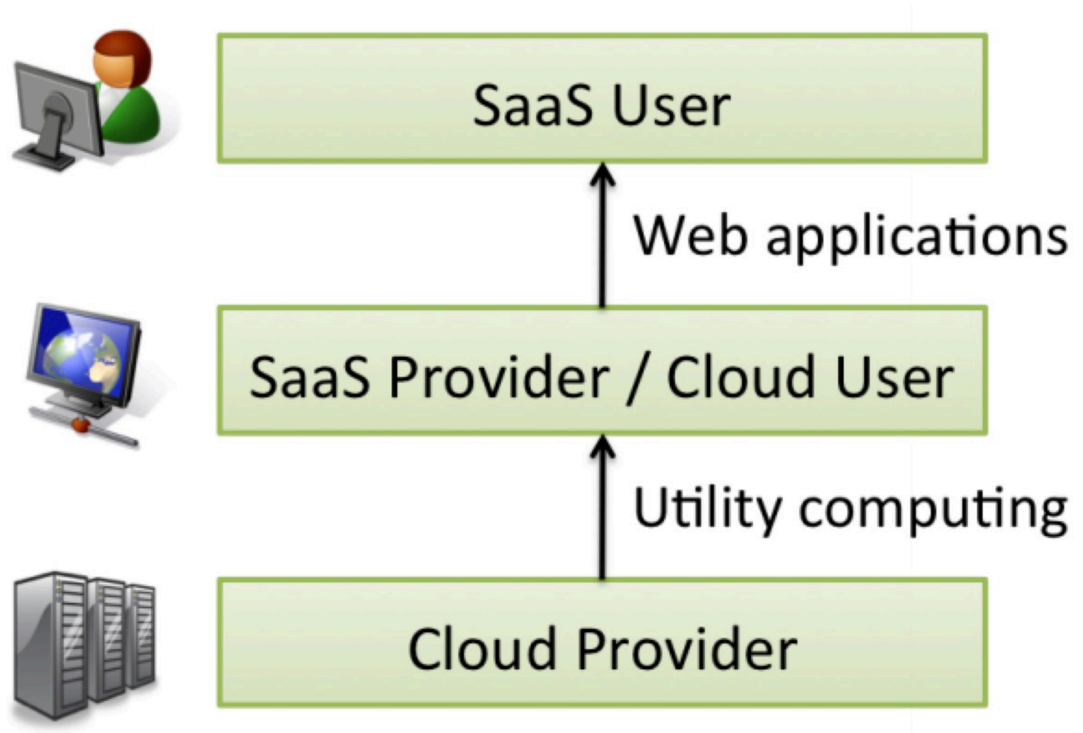
source: https://azure.microsoft.com/en-us/overview/what-is-saas/

Service	Abbr.
Analytics as a service	AnaaS
API as a service	AaaS
Artificial intelligence as a service	AIaaS
Backend as a service	BaaS
Banking as a service	BaaS
Blockchain as a service	
Business process as a service	BPaaS
Contact Information as a service	ClaaS
Content as a service	CaaS
Construction as a service	
Container as a service	
Crane as a service	
Communications as a Service	CPaaS
Data as a service	DaaS
Desktop as a service	
Drone as a service	
Database as a service	DBaaS
Distribution as a service	DaaS
Energy storage as a service	ESaaS
Electric vehicle as a service [2]	EVaaS
Function as a service	FaaS
Farming as a service	FaaS
Games as a service	GaaS
Hadoop as a service	HaaS
Housing as a service	
Infrastructure as a service	IaaS
Identity as a service	IDaaS
IT as a service	ITaaS
Logging as a service	LaaS
Management as a service	MaaS
Microgrid as a service	
Mobility as a service	
Monitoring as a service	
Metal as a service	
Mobile backend as a service	MBaaS
Machine Learning as a service	MLaaS
Network as a service	NaaS
Network Defense as a service	NDaaS
Payments as a service	PaaS
Platform as a service	
Push notification as a service	
Recovery as a service	RaaS
Robot as a service	
Search as a service	SaaS
Security as a service	
Software as a service	
Storage as a service	
Transportation as a service	TaaS
Testing as a service	TaaS
Unified Communications as a Service	UCaaS



source: https://www.skyhighnetworks.com/cloud-security-blog/microsoft-azure-closes-iaas-adoption-gap-with-amazon-aws/

Cloud Roles



Source: Above the Clouds: A Berkeley View of Cloud Computing

Cloud Models

- Who owns and runs the infrastructure?
  - Public clouds: resources shared by multiple organizations
  - Private clouds: resources used exclusively by one organization
  - Hybrid clouds

Cloud Infra Locations

- On-Premise (On-Prem)
  - Resources located locally (at a datacenter that the organization operates)
- Hosted
  - Resources hosted and managed by a third-party (cloud provider)
- Colocations (Colos)
  - A data center facility provides physical resources (e.g., building, cooling, power, etc.) for multiple organizations to host servers

- Private cloud can be on-prem, colos, and/or hosted (virtual private cloud)

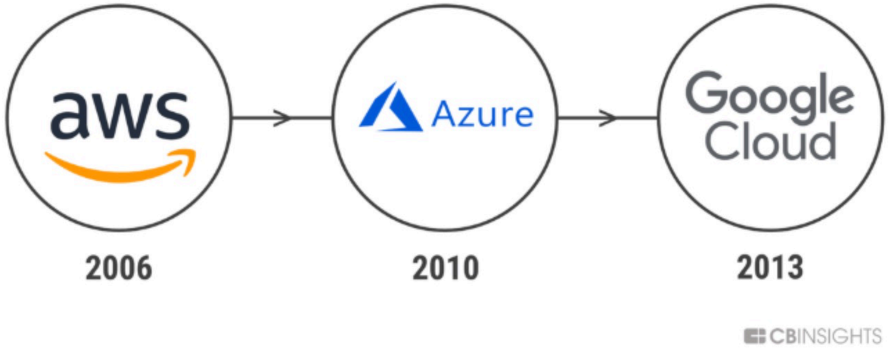
Hybrid Cloud

- Combine private (usually on-prem private) cloud and public cloud
  - Better control over sensitive data/functionalities
  - Cost effective
  - Scales well
  - Flexible

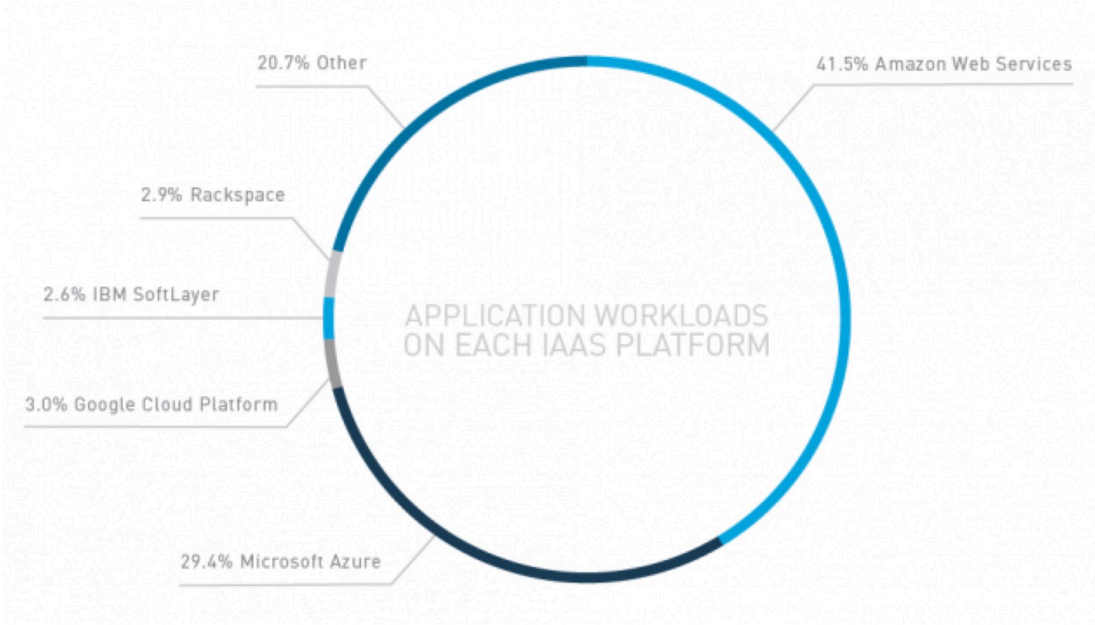
Cloud Usages

- Software/websites that serve real users (b2c)
  - Netflix, Pinterest, Instagram, Spotify, Airbnb, Lyft, Slack, Expedia
- Data analytics, machine learning, and other data services (b2b)
  - Databricks, Snowflake, GE Healthcare
- Mobile and IoT backend (b2c)
  - Snapchat, Zynga (AWS->zCloud->AWS)
- Datacenter’s own usages
  - Google Drive/OneDrive, search, internal analytics

Cloud Providers



- Companies with large datacenters, often already running large-scale software
  - Amazon AWS
  - Microsoft Azure
  - Google Cloud Platform (GCP)
  - Alibaba Cloud
  - IBM Cloud

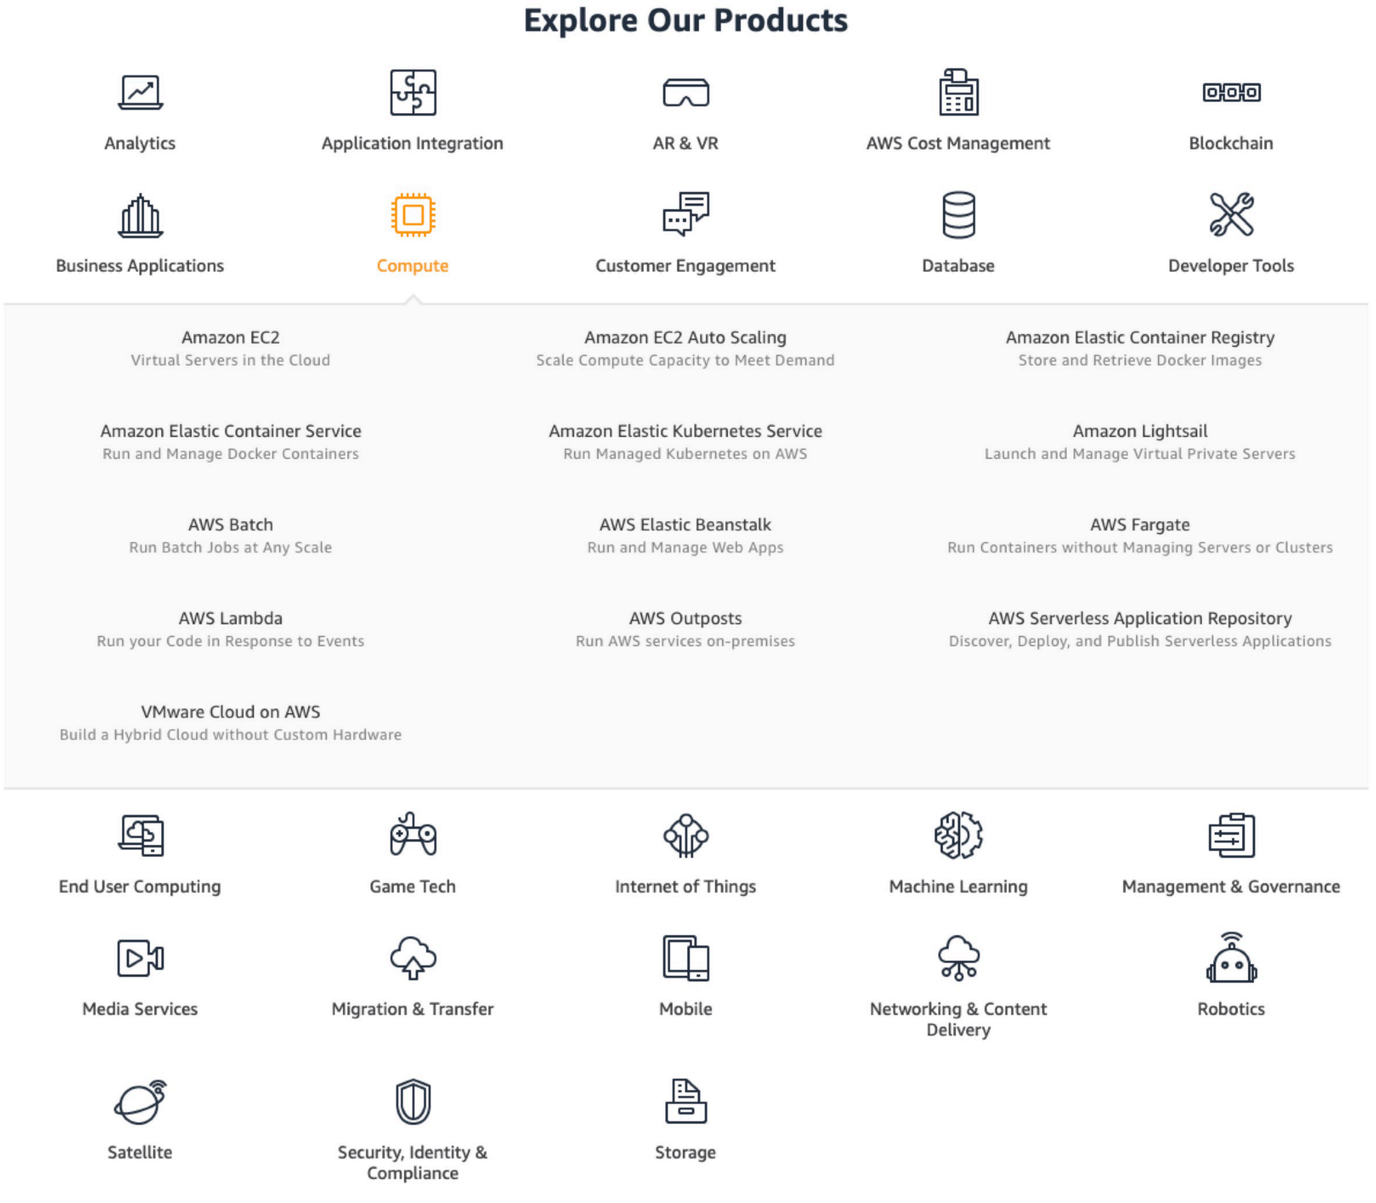


source: https://www.skyhighnetworks.com/cloud-security-blog/microsoft-azure-closes-iaas-adoption-gap-with-amazon-aws/

Amazon Web Service (AWS)

- Biggest market share, longest history
- Highest compute (and other service) options
  - >= 136 instance types in 26 families
- Storage
  - Simple Storage Service (S3)
  - Elastic Block Service (EBS)
- Many other services
  - Lambda (serverless)
  - ECS/EKS (managed containers)
  - DynamoDB, Aurora, ElastiCache (databases/key-value stores)
  - Virtual Private Cloud (VPC)
  - EMR, Redshift, many ML offerings (analytics, ML)
  - Satellite, Robotics

Cloud Services/Products/Resources



Microsoft Azure

- Moved from Windows to Linux
- Good integration with Microsoft products
  - Customers that are already using Microsoft products (e.g., having existing licenses)
- Many instance types and service types as well

Select a category:

- AI + Machine Learning
- Analytics
- Blockchain
- Compute
- Containers
- Databases
- Developer Tools
- DevOps
- Identity
- Integration
- Internet of Things
- Management and Governance
- Media
- Microsoft Azure Stack
- Migration
- Mixed Reality
- Mobile
- Networking
- Security
- Storage
- Web
- Windows Virtual Desktop

Compute

Learn more >

Access cloud compute capacity and scale on demand—and only pay for the resources you use

Virtual Machines

Provision Windows and Linux virtual machines in seconds

Service Fabric

Develop microservices and orchestrate containers on Windows or Linux

Container Instances

Easily run containers on Azure without managing servers

SQL Server on Virtual Machines

Host enterprise SQL Server apps in the cloud

SAP HANA on Azure Large Instances

Run the largest SAP HANA workloads of any hyperscale cloud provider

Virtual Machine Scale Sets

Manage and scale up to thousands of Linux and Windows virtual machines

Mobile Apps

Build and host the backend for any mobile app

Linux Virtual Machines

Provision virtual machines for Ubuntu, Red Hat, and more

Azure CycleCloud

Create, manage, operate, and optimize HPC and big compute clusters of any scale

Azure Dedicated Host PREVIEW

A dedicated physical server to host your Azure VMs for Windows and Linux

Azure Kubernetes Service (AKS)

Simplify the deployment, management, and operations of Kubernetes

App Service

Quickly create powerful cloud apps for web and mobile

Batch

Cloud-scale job scheduling and compute management

Cloud Services

Create highly-available, infinitely-scalable cloud applications and APIs

Azure Functions

Process events with serverless code

Web Apps

Quickly create and deploy mission critical web apps at scale

API Apps

Easily build and consume Cloud APIs

Windows Virtual Desktop PREVIEW

Deliver a virtual desktop experience to any device at cloud scale

Azure VMware Solution by CloudSimple

Run your VMware workloads natively on Azure

Google Cloud Platform (GCP)

- Latest among the three to come in play and smallest market share, but with good growth
- Cheapest among the three
- Fewest instance types, allows customized CPU/memory sizes
  - bill based on total CPU and memory usages, not on total instance time
- TPU support
- Good support for cross geo-regions
- More open-source projects than the other two

Google Cloud Platform

- Featured products
- AI and machine learning
- API management
- Compute
- Data analytics
- Databases
- Developer tools
- Hybrid and multi-cloud
- Internet of Things (IoT)
- Management tools
- Media and gaming
- Migration
- Networking
- Security and Identity
- Storage

More Google Cloud products

- G Suite
- Google Maps Platform
- Google Hardware
- Cloud Identity
- Chrome Enterprise
- Android Enterprise
- Apigee
- Firebase
- Orbitera
- Product launch stages

Compute →

Compute Engine
Scalable, high-performance VMs.

Shielded VMs
Hardened virtual machines on GCP.

Google Kubernetes Engine (GKE)
Run containerized applications.

Anthos on-premises
Build and manage modern hybrid applications on existing VMware environments.

Container security
Secure your container environment on GCP.

Migrate for Compute Engine
Purpose-built, enterprise-grade migration to Google Cloud.

App Engine
Serverless application platform for apps and backends.

Cloud Run (beta)
Run stateless containers on a fully managed environment or on Anthos.

Cloud Functions
Event-driven serverless compute platform.

Cloud Functions for Firebase
Run mobile backend code without managing servers.

Knative
Components to create modern, Kubernetes-native cloud-based software.

Graphics Processing Unit (GPU)
Leverage GPUs on Google Cloud for machine learning, scientific computing, and 3D visualization.

Data analytics →

BigQuery
A fully managed, highly scalable data warehouse with built-in ML.

Cloud Dataflow
Real-time batch and stream data processing.

Cloud Dataproc
Managed Spark and Hadoop service.

Cloud Datalab
Explore, analyze, and visualize large datasets.

Cloud Dataprep

Cloud Composer
A fully managed workflow orchestration service built on Apache Airflow.

Cloud Data Fusion
Fully managed, code-free data integration.

Data Catalog
A fully managed and highly scalable data discovery and metadata management service.

Genomics
Power your science with Google Genomics.

Multi-Cloud

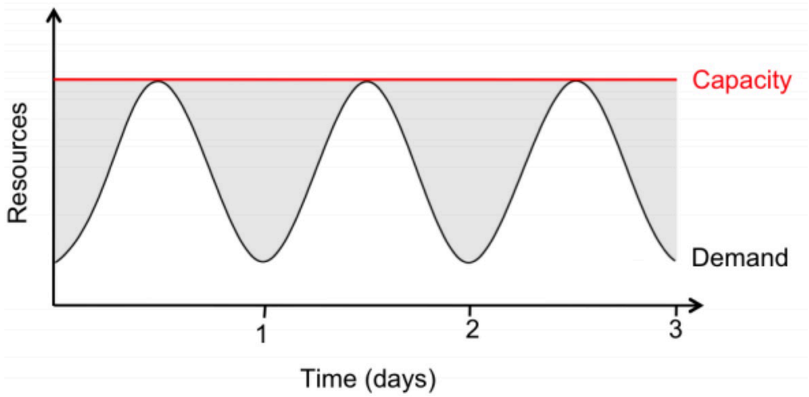
- Use multiple clouds for an application/service

- Avoid data lock-in
- Avoid single point of failure
- Need to deal with API differences and handle migration across clouds

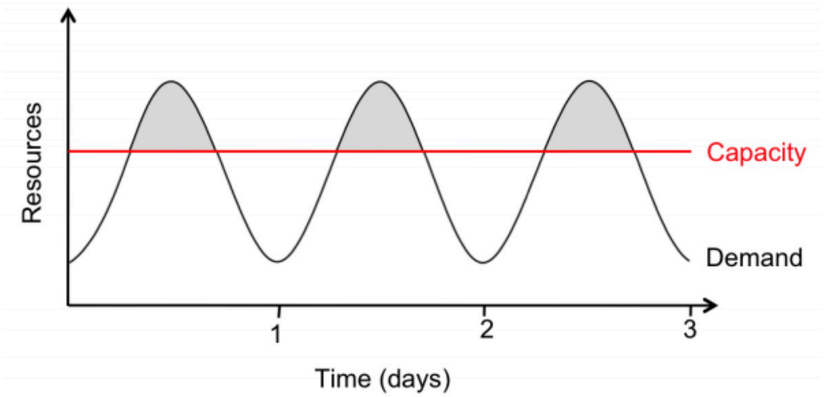
Outline

- What is cloud computing?
- Why cloud computing?
- How clouds are built?

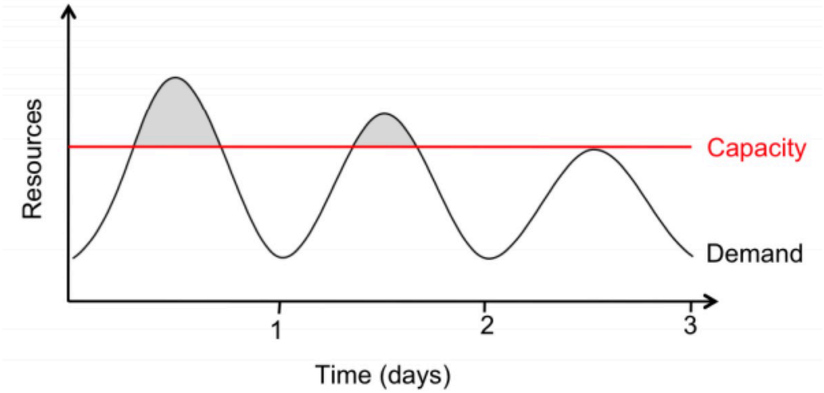
Incentive for Cloud Users



(a) Provisioning for peak load



(b) Underprovisioning 1



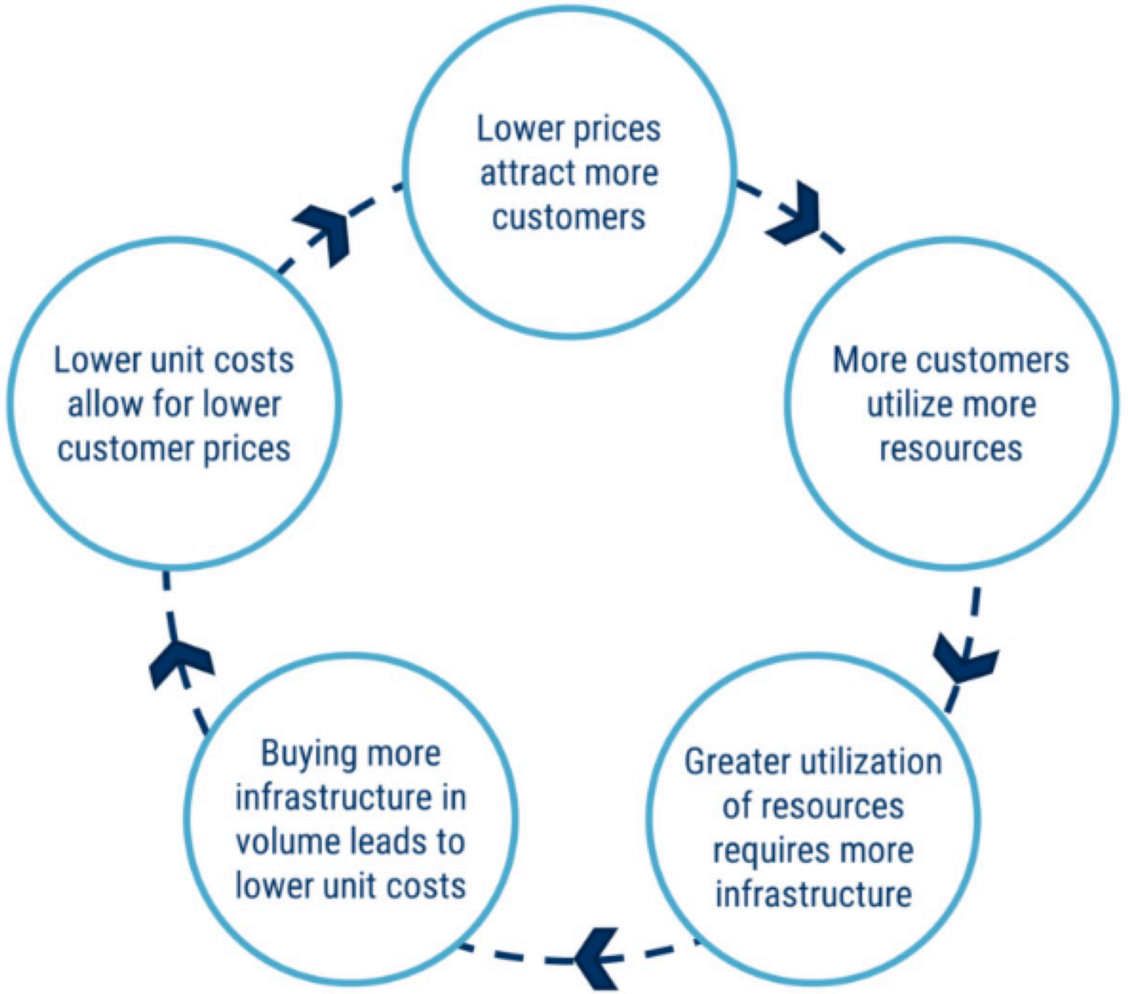
(c) Underprovisioning 2

Source: Above the Clouds: A Berkeley View of Cloud Computing

Incentive for Cloud Providers

- Make a lot of money
- Leverage existing investment
- Defend a franchise
- Attack an incumbent
- Leverage customer relationships
- Become a platform

The Virtuous Cycle of Cloud Computing



Source: CloudBlogs.Microsoft.com

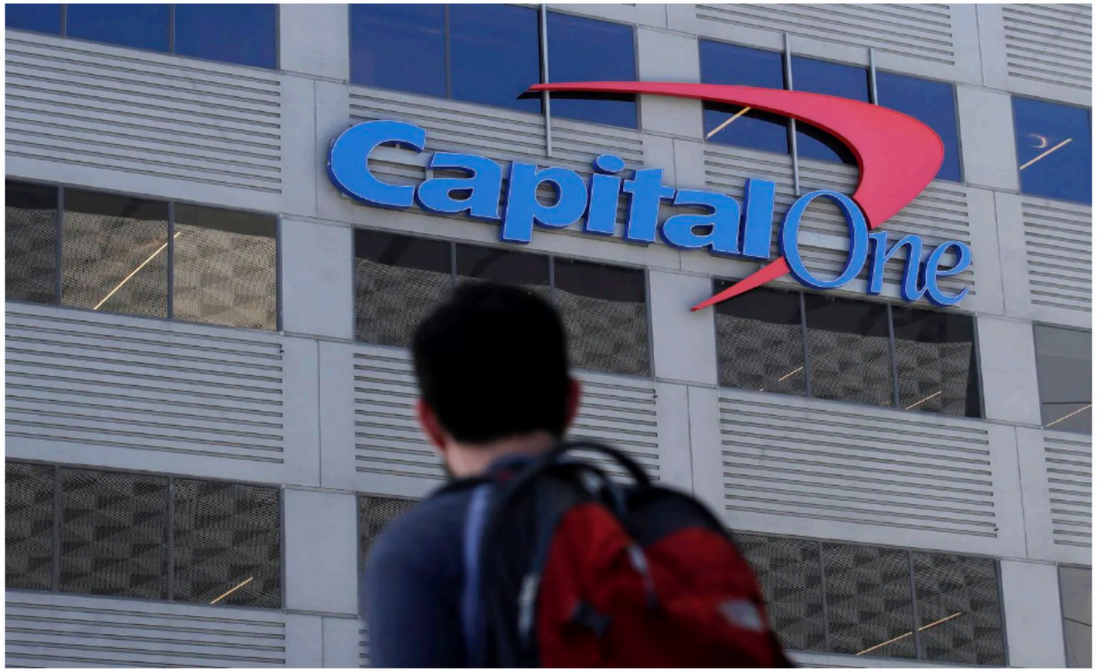
CBINSIGHTS

The Washington Post
Democracy Dies in Darkness

Technology

The Capital One hack couldn't have come at a worse time for Amazon's most profitable business

Amazon's cloud-computing unit was considered the front-runner in a race to secure a $10 billion Pentagon contract. Now that could be in trouble.

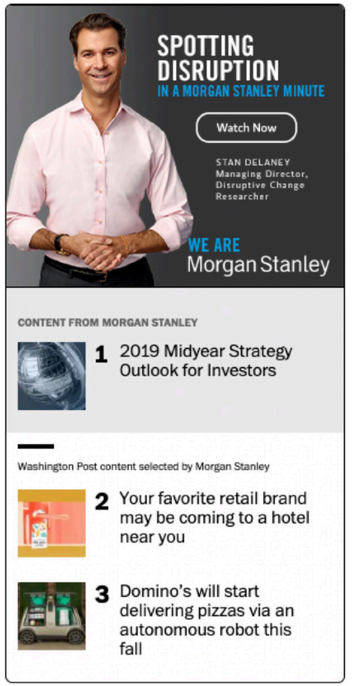


A Capital One location in San Francisco. (Jeff Chiu/AP)

By **Jay Greene** and **Drew Harwell**
August 1

The data breach that appears to have exposed more than 100 million applications for Capital One credit cards couldn't have come at a worse time for Amazon Web Services, which stores the bank's data.

The profit-driving Amazon unit, which allows companies to rent out storage and computing space on massive servers, has been the favorite to win a 10-year, $10 billion contract from the Defense Department, which had been expected to be announced this month. It was thrown into question



How the Jennifer Lawrence iCloud hack really happened



Chris Smith @chris_writes
March 16th, 2016 at 6:50 AM

Share Tweet

In early September 2014, Apple was preparing to announce a massive iPhone upgrade: bigger iPhones than anything it launched before. But just a few days ahead of the press iPhone 6 press event, an iCloud security scandal broke out. Nude pictures belonging to Jennifer Lawrence and many other celebrities leaked online, originating from iPhone backups.

Apple explained at the time that its iCloud security was not breached and that hackers probably employed phishing schemes to obtain the usernames and passwords from their victims.

Now, 18 months later after the scandal, we finally find out what happened. And it turns out that phishing attacks were indeed used to target the celebrities.

DON'T MISS: 11 paid iPhone apps on sale for free for a limited time

According to *NBC News*, it's 36-year-old Ryan Collins the person responsible for phishing login credentials from many celebrities. With usernames and passwords in hand, he was able to log into Gmail accounts and even download iCloud backups from where he extracted nude photos.

Source: https://bgr.com/2016/03/16/jennifer-lawrence-nudes-icloud-hack/

alas.aws.amazon.com/ALAS-2019-1222.html

Amazon Linux Security Center

ALAS-2019-1222

Amazon Linux AMI Security Advisory: ALAS-2019-1222
Advisory Release Date: 2019-06-17 17:58 Pacific
Severity: ⚠ Critical
References: CVE-2019-11479 CVE-2019-11478 CVE-2019-11477

Issue Overview:
CVE-2019-11477 □, CVE-2019-11478 □ and CVE-2019-11479 □ describe vulnerabilities in the Linux kernel that can be remotely exploited using a specially crafted TCP connection, crashing the targeted system.

The latest Amazon Linux AMIs as available in AWS EC2 already contain these kernels and are not vulnerable.

Affected Packages:
kernel

Issue Correction:
Run *yum update kernel* and reboot to update your system.



The Impact of GDPR on Cloud Computing

22 May 2018

The EU General Data Protection Regulation (GDPR) is the most significant piece of privacy legislation to come into effect across Europe in a generation.

It will apply to any organisation who handles EU citizen data, even if they're not from the EU. This means any European company with employees or a US firm with European customers will have to comply or face the consequences.

GDPR places new obligations on businesses, ones which will affect how they use cloud services. With cloud adoption now around 90% in the UK, it is important to ensure that the cloud services you use are compliant and that the systems and applications you design do not expose you to risk.

Strengthening privacy

GDPR strengthens user privacy in two main ways. Firstly, it increases the obligations on organisations to protect user data and secondly it grants citizens major new powers over how their information is collected, used and stored.

For example, businesses must ensure that all reasonable steps are taken to secure data, train staff and disclose breaches and must be clear and transparent to citizens about how they use personal data. Citizens can demand to see what data an entity is held on them and can also request that this is deleted at any time.

Source: https://www.uk.insight.com/en-gb/content-and-resources/articles/2018-05-22-the-impact-of-gdpr-on-cloud-computing

Obstacles and Opportunities

	Obstacle	Opportunity
1	Availability of Service	Use Multiple Cloud Providers; Use Elasticity to Prevent DDOS
2	Data Lock-In	Standardize APIs; Compatible SW to enable Surge Computing
3	Data Confidentiality and Auditability	Deploy Encryption, VLANs, Firewalls; Geographical Data Storage
4	Data Transfer Bottlenecks	FedExing Disks; Data Backup/Archival; Higher BW Switches
5	Performance Unpredictability	Improved VM Support; Flash Memory; Gang Schedule VMs
6	Scalable Storage	Invent Scalable Store
7	Bugs in Large Distributed Systems	Invent Debugger that relies on Distributed VMs
8	Scaling Quickly	Invent Auto-Scaler that relies on ML; Snapshots for Conservation
9	Reputation Fate Sharing	Offer reputation-guarding services like those for email
10	Software Licensing	Pay-for-use licenses; Bulk use sales

Source: Above the Clouds: A Berkeley View of Cloud Computing

10 years later

- What’s your view on Cloud Computing?

Outline

- What is cloud computing?
- Why cloud computing?
- How clouds are built?

Virtualization

- Traditional: applications run on physical servers
  - Manual mapping of apps to servers
    - Apps can be distributed
    - Storage may be on a SAN or NAS
  - IT admins deal with “change”

- Modern: virtualized data centers
  - App run inside virtual servers; VM mapped onto physical servers
  - Provides flexibility in mapping from virtual to physical resources

Virtualization Benefit

- Resource management is simplified
  - Application can be started from preconfigured VM/container images
  - Virtualization layer / hypervisor permits resource allocations to be varied dynamically
  - VMs can be migrated without application down-time

Virtual Datacenter

- A cluster of machines, each running a set of VMs
  - drive up utilization by packing many VMs onto each cluster node
  - fault recovery is simplified
    - if hardware fails, copy VM image elsewhere
    - if software fails, restart VM from snapshot
  - can safely allow third parties to inject VM images into your data center
    - hosted VMs in the cloud, commercial computing grids

Recent Trend: Container

- Light-weight virtualization
  - Running multiple isolated user-space applications on one OS
  - Virtualization layer runs as an application within the OS
  - Focusing on resource and namespace isolation

- Example: Docker, LXC, Kubernetes, Xen Unikernel

Software-Defined Data Center

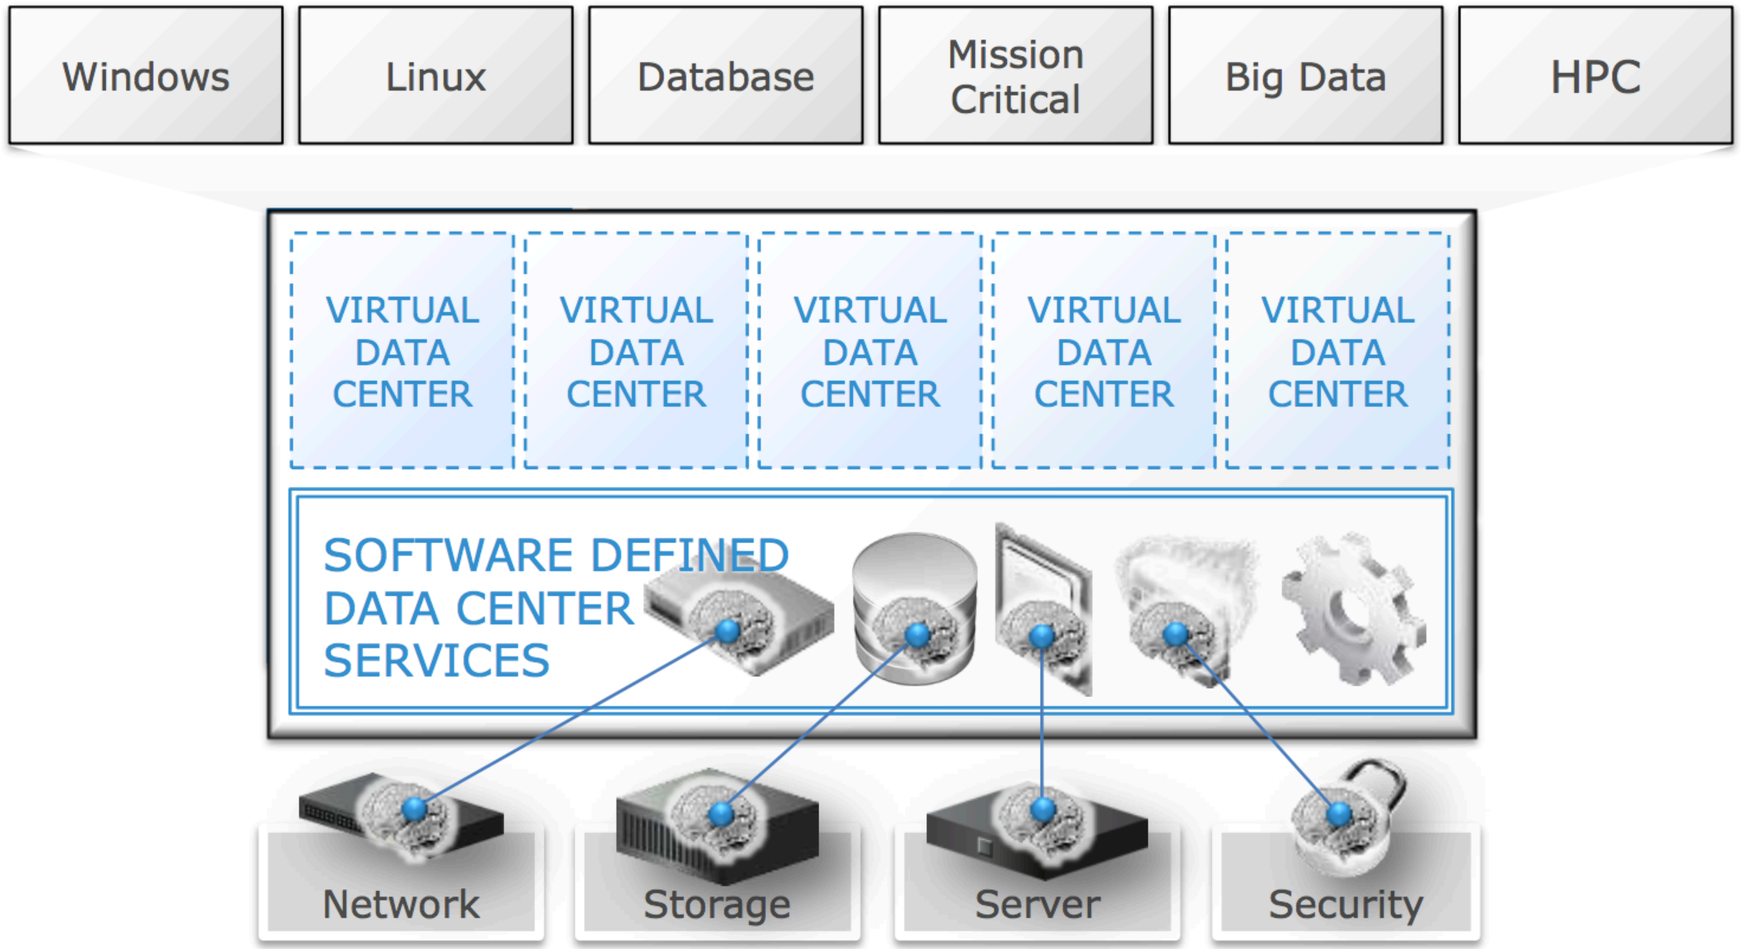
- All infrastructure is virtualized and delivered as a service & the control of this datacenter is entirely automated by software

Traditional Data Center



Source: EMC

Software-Defined Data Center



Source: EMC