Cloud Computing Use Cases

15-719 Advanced Cloud Computing

Garth Gibson
Greg Ganger
Majd Sakr

Adoption of virtual desktops by government agencies

Garth Gibson, Greg Ganger, Majd Sakr

15-719/18-8467b: Advanced cloud computing
Spring 2017

Adopted from Raja Sambasivan, Fall 2013

Cloud-based clients (i.e., virtual desktops)

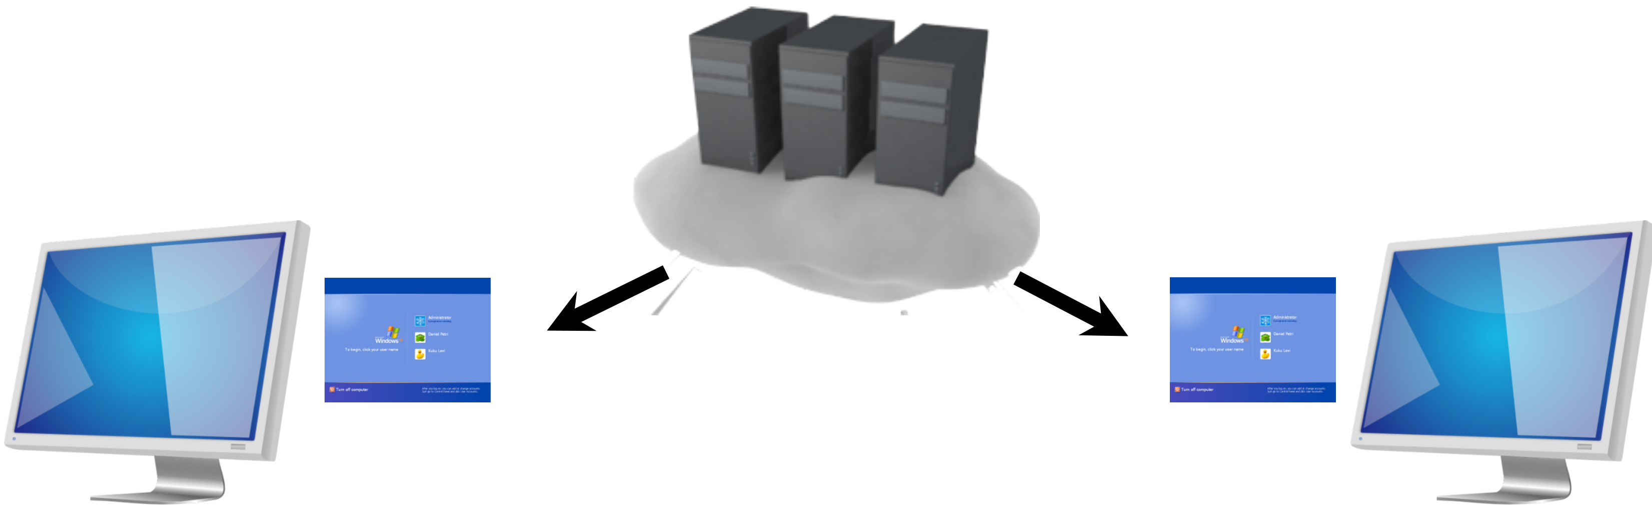

- Cloud stores users' desktop images
- Users use desktop remotely via thin clients
  - Most computation done on cloud

Clients in traditional IT environments

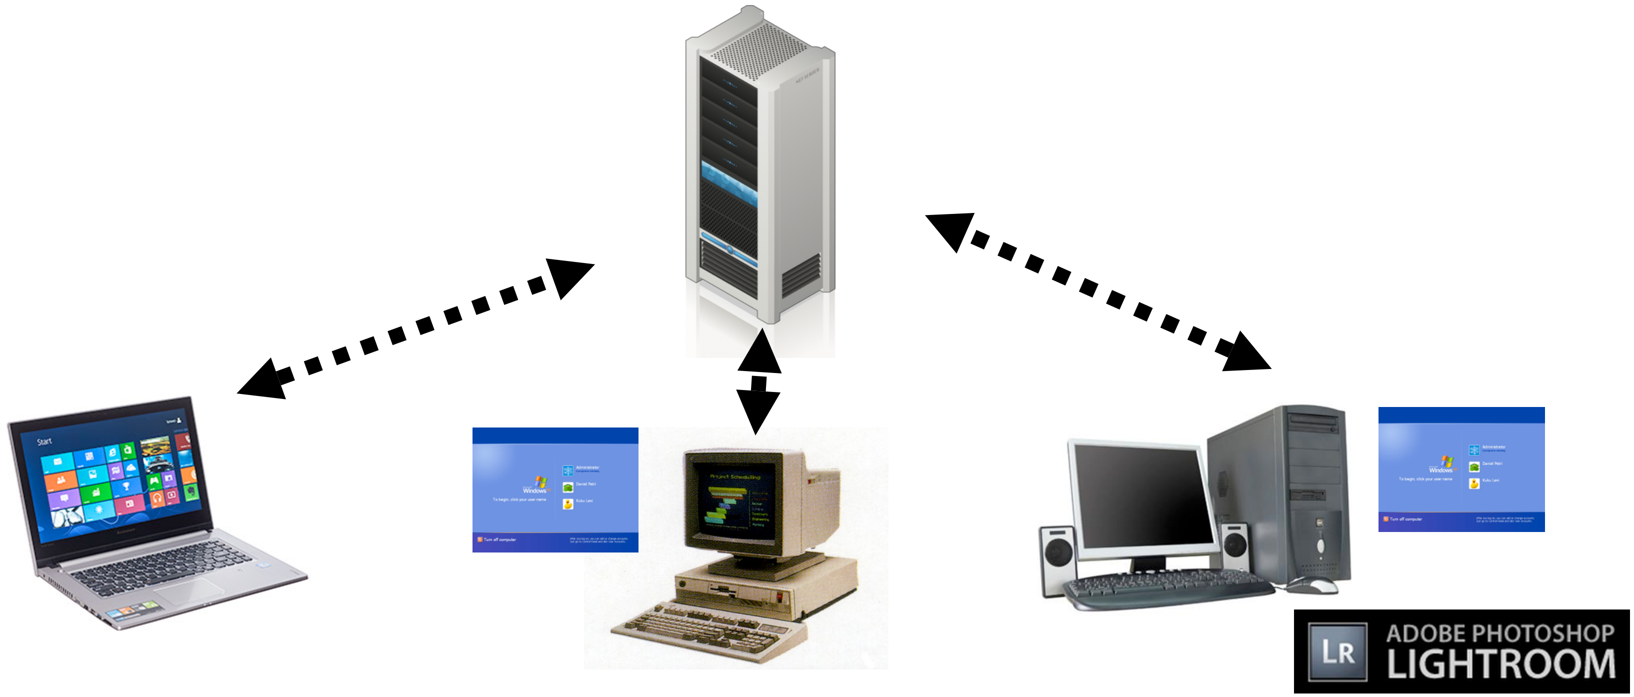

- Machines run locally-installed software
- Software/machine configurations can vary

Advantages of cloud-based clients (I)

- Reduces admin/support cost & effort
  - Updates/patches can be applied centrally
    - No need to rely on users (or travel to user sites)
  - Eases backup/recovery tasks
  - More homogeneity makes support easier
  - User-specific hardware has less impact

Advantages of cloud-based clients (II)

- Reduces fixed cost of per-user hardware
  - Enables BYOD, no need to upgrade as often
- Increases organizational security
  - Access to lost/stolen devices can be disabled
- Allows users to work from anywhere

City Information Systems (Pittsburgh, PA)

- Supports 2,800 users (City Hall, Police, etc.)
- Moved to virtual desktops to up efficiency
  - Necessary due to 2008 budget cuts
- **Results**:
  - Can handle 10 times more calls/day
  - PC retirement age gone up from 5 to 7 years

Source: http://www.vmware.com/files/pdf/resources/IDC_Whitepaper234253.pdf

Challenges to adoption

- Cost of new infrastructure needed
  - E.g., for more storage/network bandwidth
- Last-mile network bandwidth
- Security (specifically for public cloud)

Drawbacks

- Can't work offline
- Central point of failure
- Less user customization/expressivenesss

Advanced Cloud Computing Use Cases

- Reference 1: "Characterizing Private Clouds: A Large-Scale Empirical Analysis," Cano, Aiyar, Krishnamurthy, SoCC16
- Opt Reference 2: "Why Virtual Desktop at CCRI? Finding sustainability for desktop support," Vieira, SIGUCCS12.
- Opt Reference 3: "Desktop virtualization in state and local governments: saving money wothout sacrificing citizen service," Clarke, IDC gov't insights (sponsored by VMWare), 2012.
- Opt. Reference 4: "AWS Summit New York 2016 Keynote," Werner Vogels (Amazon CTO), 2016.
  - Other: AWS Case Study: Obama for America
  - Other: Virtual Supercomputing with CycleCloud

Obama for America (OFA): Re-election campaign

- Use rapid data integration and predictive analytics to win elections
  - Randomized experiments (data mined) to target and influence voters in states where the races were close
- Data sources: demographic data, voting history data, fundraising data, volunteer data, social media data, polling data
- Channels: door knocks, phone calls, direct mail, email, online ads, social media, paid media/TV, web
- Goals: recruit volunteers, raise funds, register voters, persuade voters, "get out the vote (GOTV)" to the polls

Short-lived intensely scaled project (& corporation)

- tweet: by OFA head of DevOps:
  - 4 Gbps
  - 10,000 requests per second
  - 2,000 nodes
  - 3 datacenters
  - 180 TB
  - 8.5 billion requests
  - Designed, deployed and dismantled in 583 days
- Admits to only 30 mins of downtime
- Core tech team of 40 people

Obama for America (OFA): Re-election campaign

- Built 200 applications, with constant change in computing environment, including …
- Narwhal – REST API (python) for common data repository for all services
  - Most data in MySQL-based Relational Database Service (RDS) but also used:
    PostgresDB, MS SQL Server, MongoDB, Vertica, LevelDB, S3, DynamoDB, SimpleDB
  - Integration and triggering driven by Simple Queue Service (SQS)
  - Based on Service Oriented Architecture (SOA), a way of organizing modular services for reuse
- CallTool – 1000s of volunteers calling millions of voters over last four days of the campaign, matching volunteer to voters and done from volunteer homes
- Dashboard – online (Rails) tool for registering & tracking volunteers and information they gather
- Dreamcatcher – processed social network sites for political sentiments to micro-target voters that could be won over
- Identity – single signon for all services; statistics and permission for users
- Ushahidi – voter incident tracking tool taken from humanitarian open source
- GOTV – Get Out The Vote application used social networking (& DynamoDB) to get voters to polls

AWS-OFA Architecture

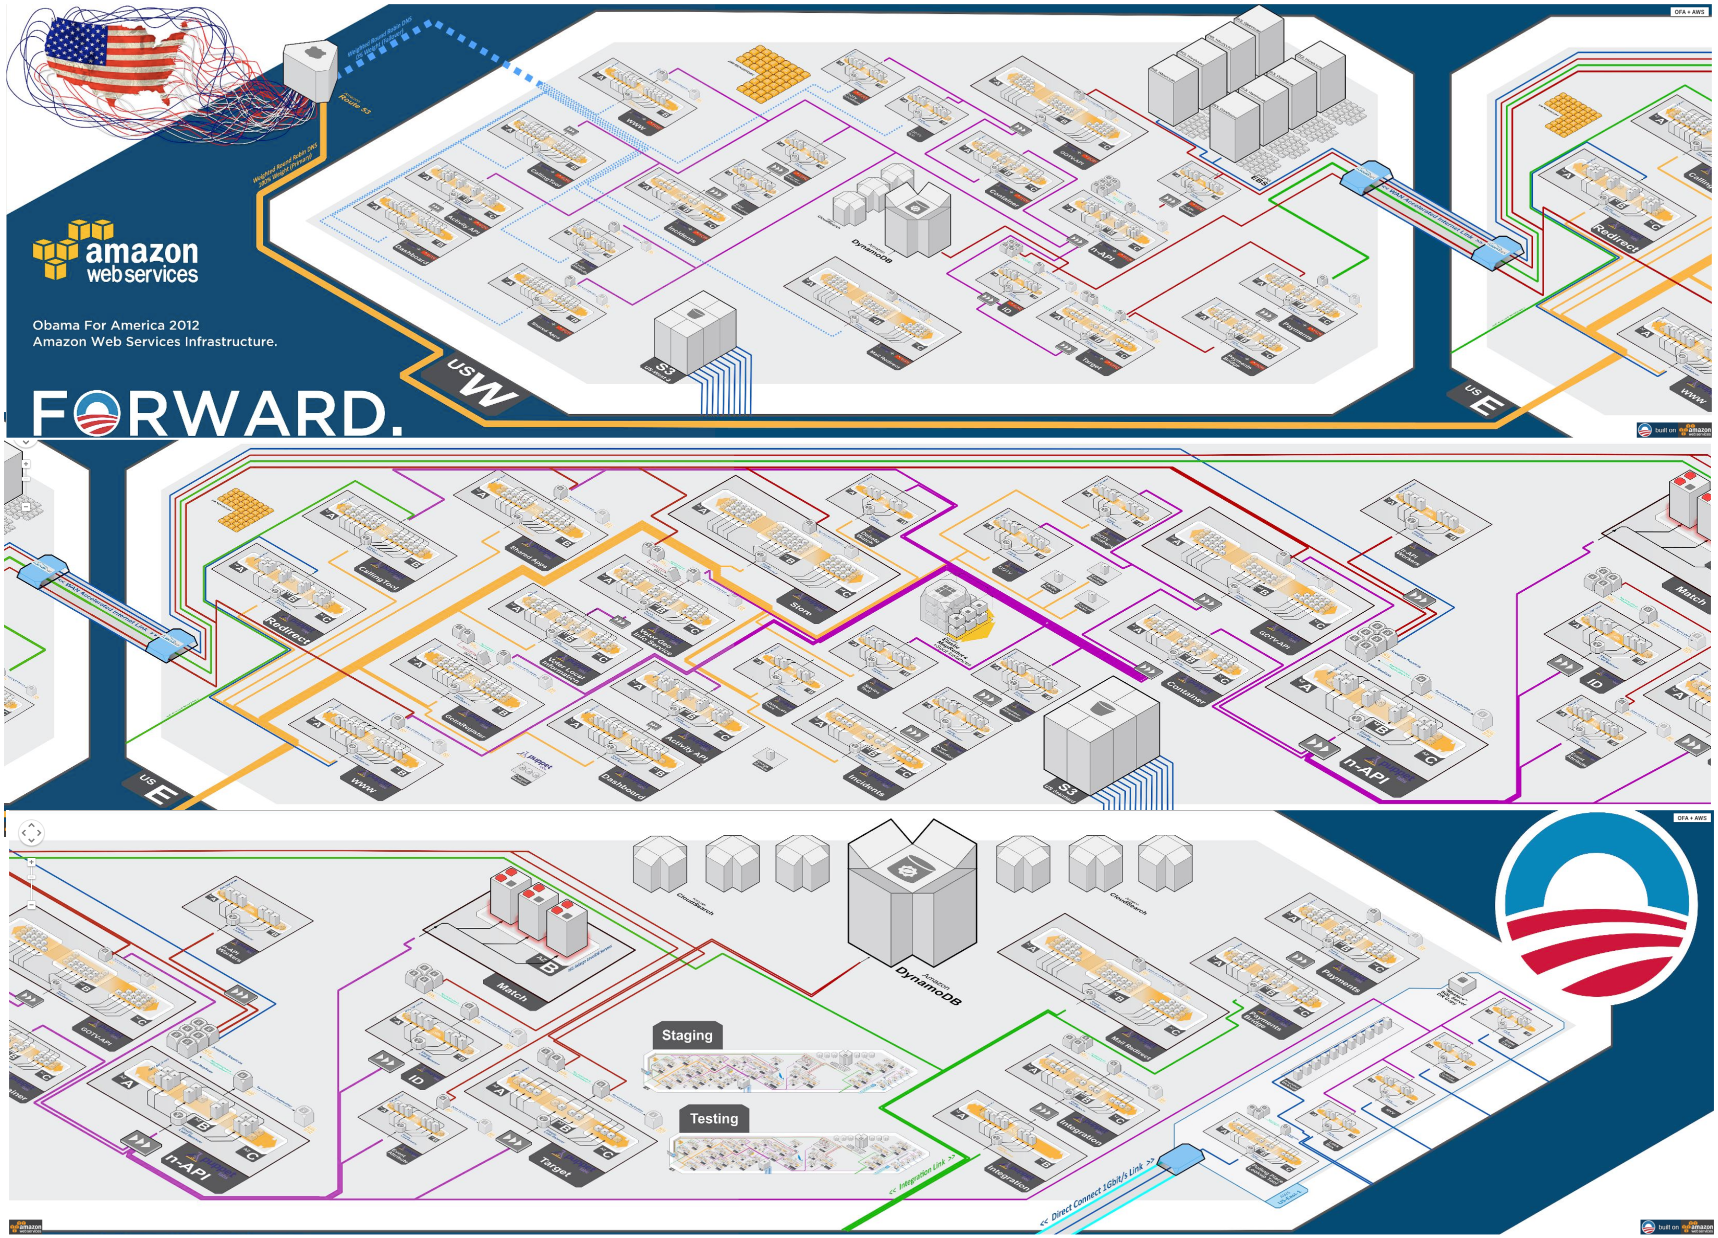

Lots of replication, backups and snapshots

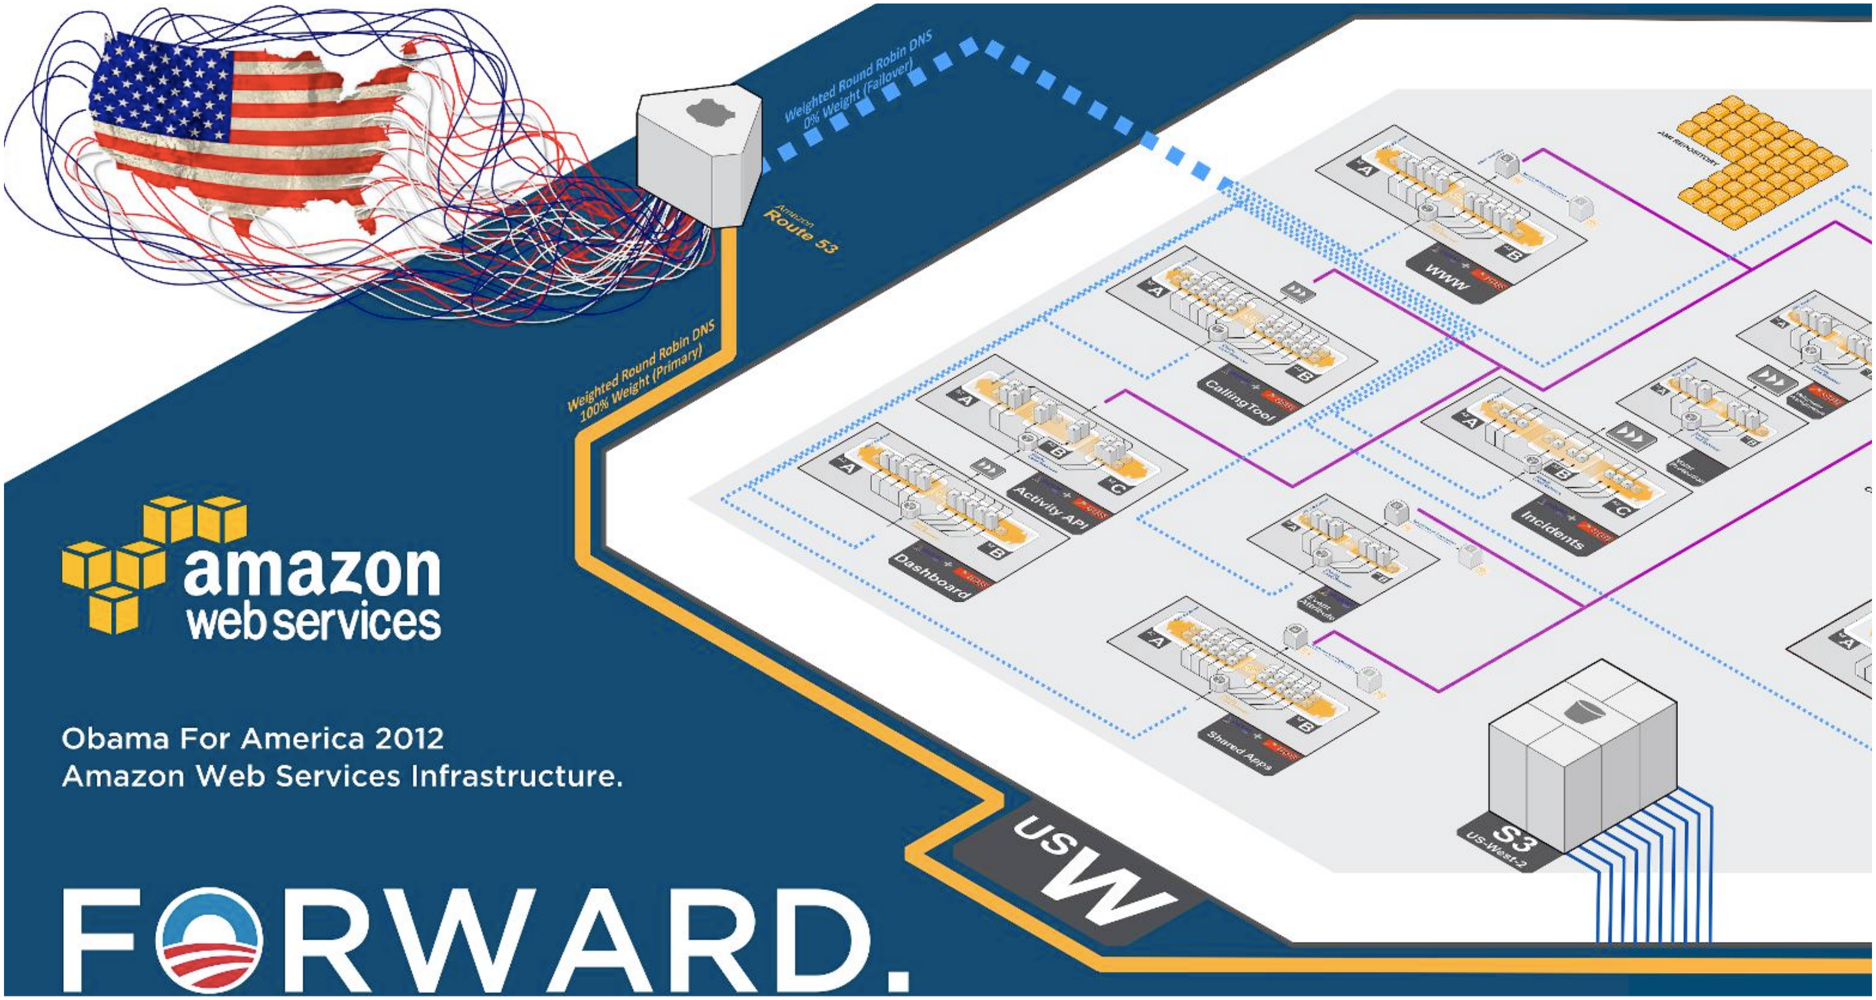

Services are load balanced clusters + replicated RDS

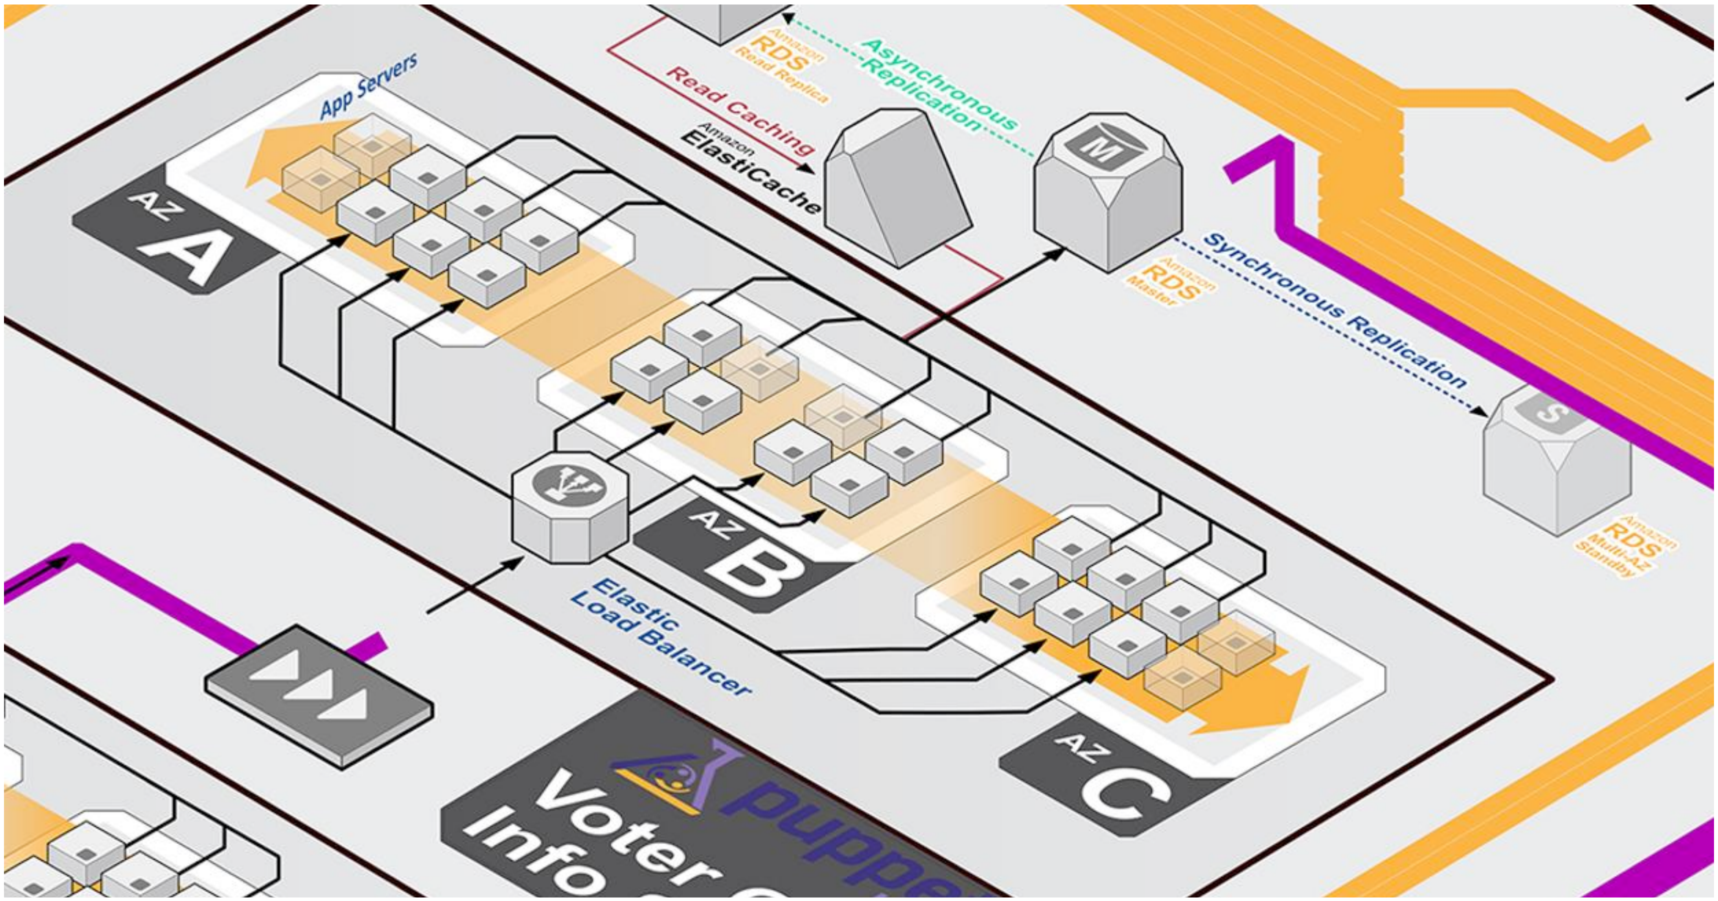

Amazon Cloud Management Tools

- AutoScaling with Elastic Load Balancing (based on CloudWatch)
  - With elasticity support from Puppet, CloudFormation and Asgard
- Multi-AZ with RDS – synchronous replication in another availability zone
- SSH and AWS Access Key rotation policies
- Route53 DNS management – scalable, available DNS implementation
- Access control based on security group identities (IAM)
- Virtual Private Clouds (VPC) – space sharing, user managed VPC network
- Direct Connect – dedicated network connections from offsite locations and integrated VLANs

Note isolated copies for testing, and staging changes

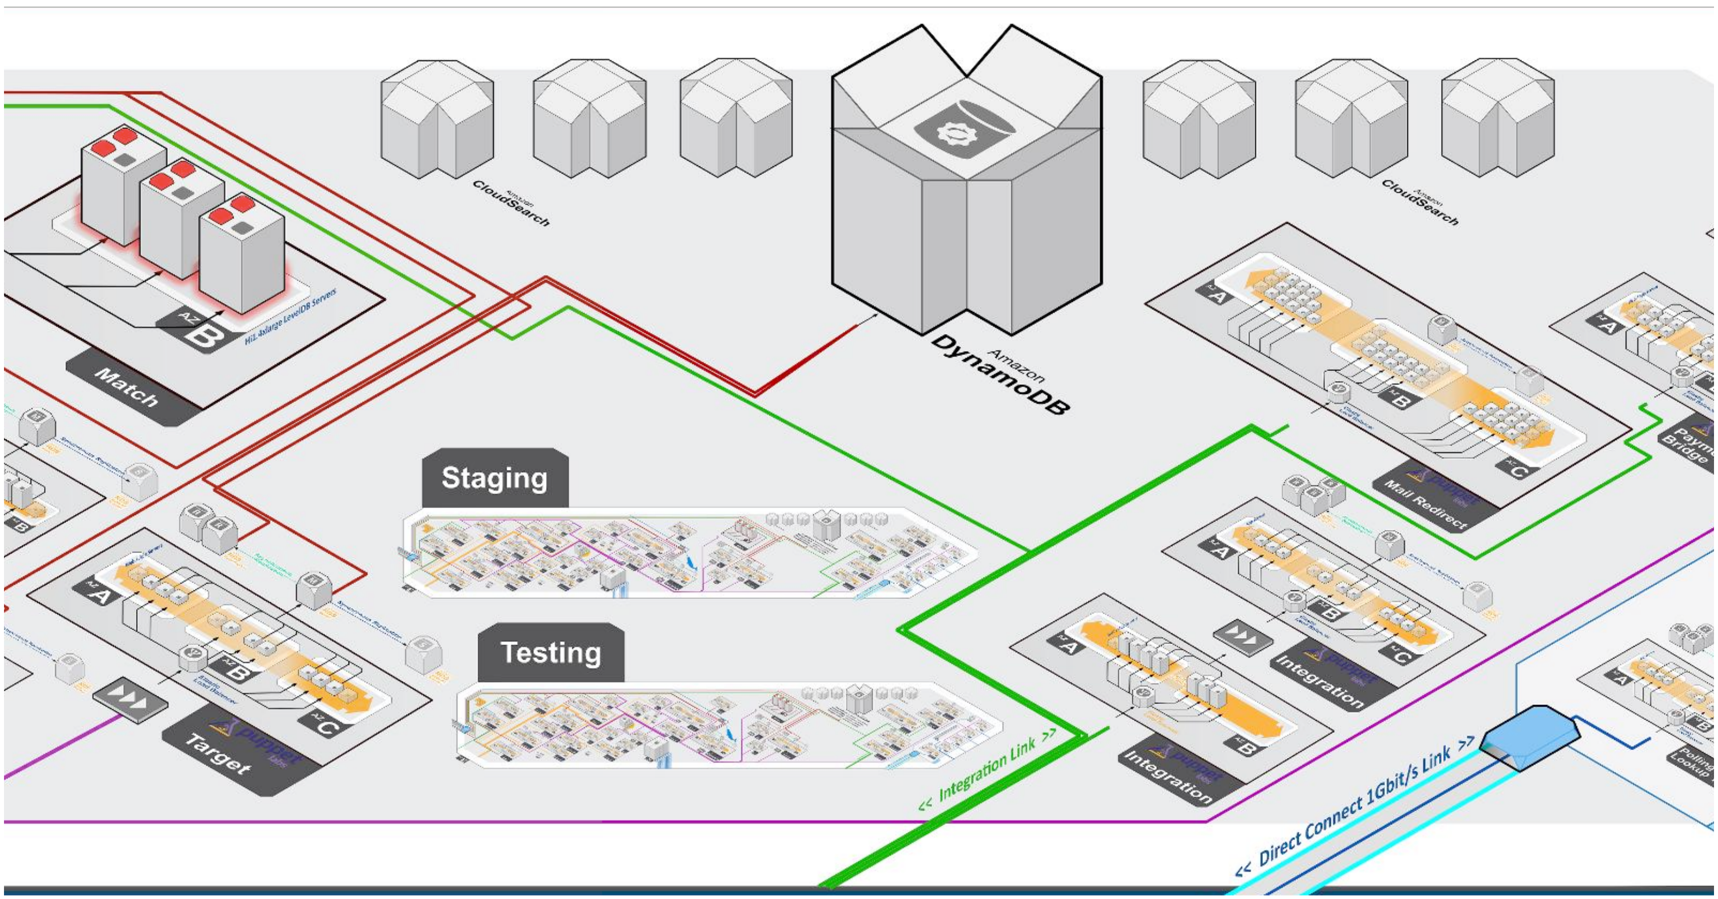

Cycle Computing: On-demand Supercomputers

- Business eg.: Varian wanted to design a mass spectrometer, needing 6 weeks simulation on the inhouse cluster, but didn't have that time
- Cycle Computing use case drove development of CycleCloud, which is a value-add broker for cloud vendors
- Werner quote: "Internal HPC clusters are too small when you need them most …. And too large every other time."

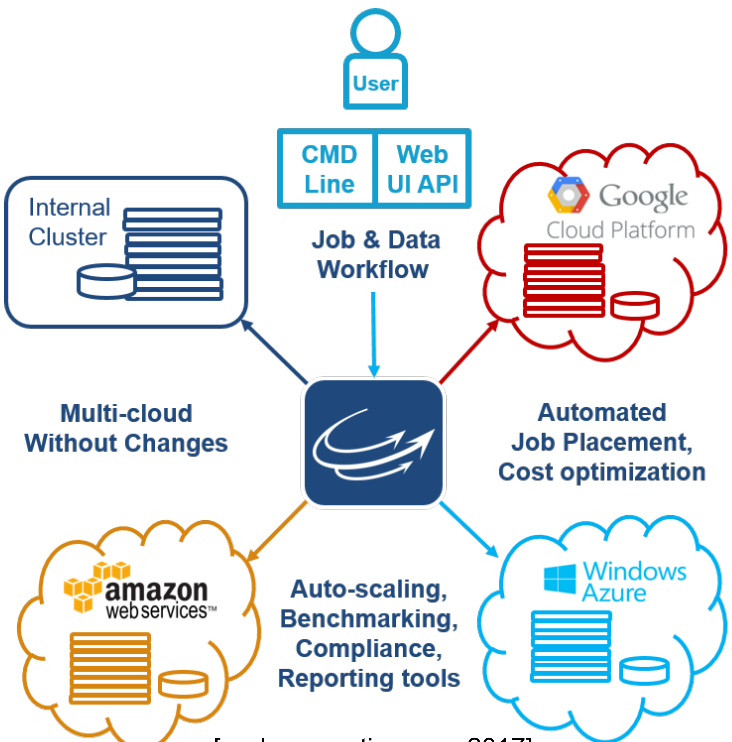

[cyclecomputing.com 2017]

CycleCloud Virtual HPC Cluster

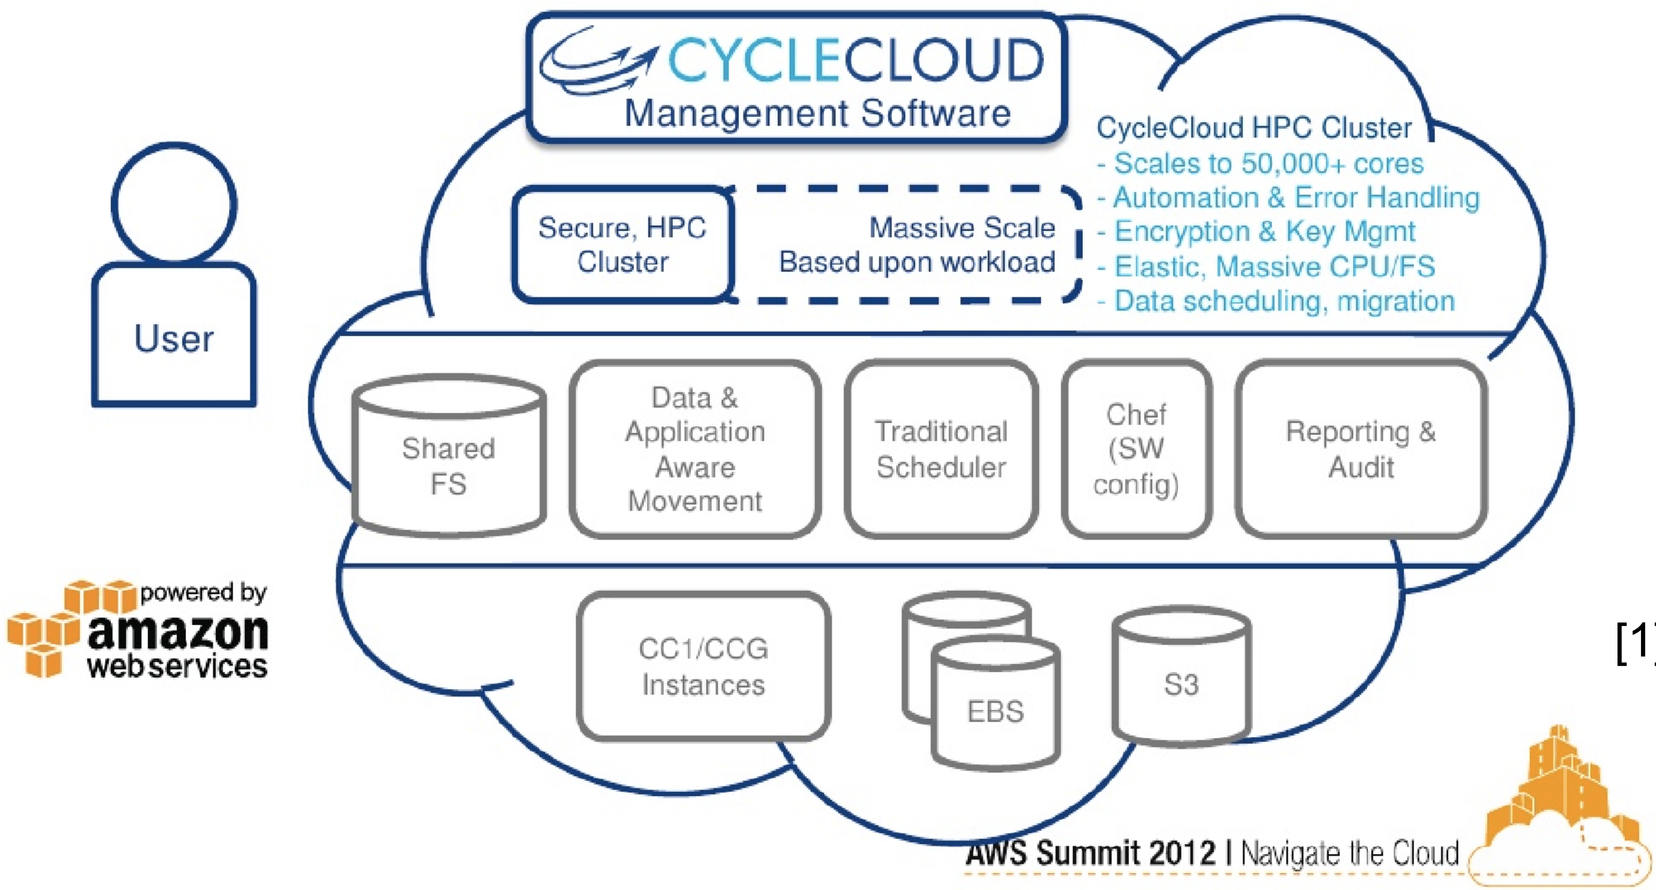

[1]

Packaged Apps

- Bioinformatics
  - Genome search
  - Gene alignment
- Proteomics
  - 3D protein structure
- Computational Chemistry
  - Simulation of fluid dynamics and molecular interactions

Bioinformatics

Blast	BLAST (Basic Local Alignment Search Tool) is an algorithm for comparing primary biological sequence information, such as the amino-acid sequences of different proteins or the nucleotides of DNA sequences.
GMAP	A genomic mapping and alignment program for mRNA and EST sequences.
HMMER	HMMER is used for searching sequence databases for homologs of protein sequences, and for making protein sequence alignments. It implements methods using probabilistic models called "profile hidden Markov models".
MAQ	Maq stands for Mapping and Assembly with Quality It builds assembly by mapping short reads to reference sequences.
Bowtie	Bowtie is an ultrafast, memory-efficient short read aligner.
RMAP	RMAP is aimed to accurately map reads from the next-generation sequencing technology. RMAP can map reads with or without error probability information (quality scores) and supports paired-end reads or bisulfite-treated reads mapping.
MrBayes	MrBayes is a program for the Bayesian estimation of phylogeny.

Proteomics

OMSSA	The Open Mass Spectrometry Search Algorithm [OMSSA] is an efficient search engine for identifying MS/MS peptide spectra by searching libraries of known protein sequences.
X! Tandem	X! Tandem open source is software that can match tandem mass spectra with peptide sequences, in a process that has come to be known as protein identification.

Computational Chemistry

GROMACS	GROMACS is a versatile package to perform molecular dynamics, i.e. simulate the Newtonian equations of motion for systems with hundreds to millions of particles.
ROCS	ROCS is a fast shape comparison application, based on the idea that molecules have similar shape if their volumes overlay well and any volume mismatch is a measure of dissimilarity.

15-719/18-847b
Advanced Cloud Computing

Lecture 02

Cloud Use Cases – Start-ups

January 23, 2017

http://www.cs.cmu.edu/~15719/

Computing Needs of a Tech Start-up

- Heterogeneous computing needs
  - Engineering, software development, business development, hosting, etc.
- Low volume
- Rapidly changing
- Low cost (when possible)

On Premise System Challenges

- Upfront capital investment
- Lengthy procurement cycles
- Lengthy deployment effort
- Costly power and cooling
- Costly system administration
- Low utilization
- Costly disaster recovery

Start-up Cloud Adoption

- Save capital expenditure which could be used to drive other areas of business growth
- Heterogeneous infrastructure on demand
- Save money on data-center real estate, personnel, power and cooling
- Scale infrastructure as the business grows
- Levels the infrastructure playing field with established companies

GoSquared

- Real-time web analytics
  - Free trial for everyone
    - John Smith
    - Wall Street Journal
  - Unpredictability with who signs up
    - Needing to scale immediately
- Covered on TechCrunch
  - Not enough resources to respond to load
- Needing to re-architect frequently
  - Rapidly redesign and redeploy their solution

15-719/18-847b
Advanced Cloud Computing

Lecture 02

Cloud Use Cases – Data Analytics

January 23, 2017

http://www.cs.cmu.edu/~15719/

Evidence Based Decision Making

- Data Science driven by data analytics
- Many different sources of multi-media data
  - Text, images, audio, video
  - Static and real-time
- Batch processing of static data
  - Non-iterative processing
  - Iterative processing
- Stream processing of real-time data

Life Cycle of Data Analytics

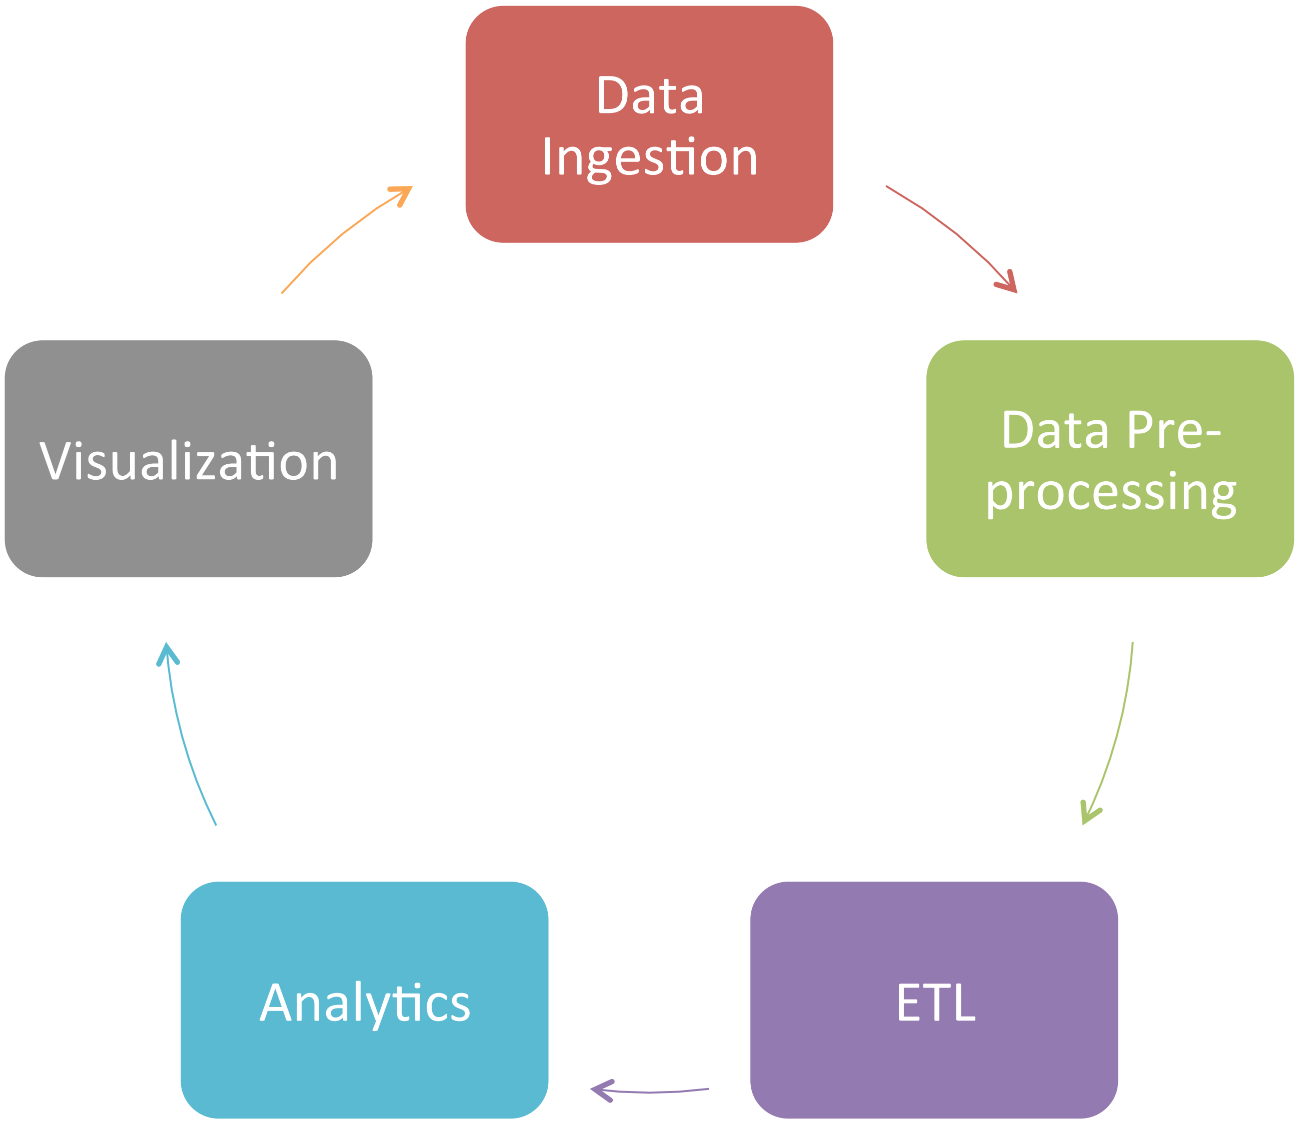

Cloud Services for Data Analytics

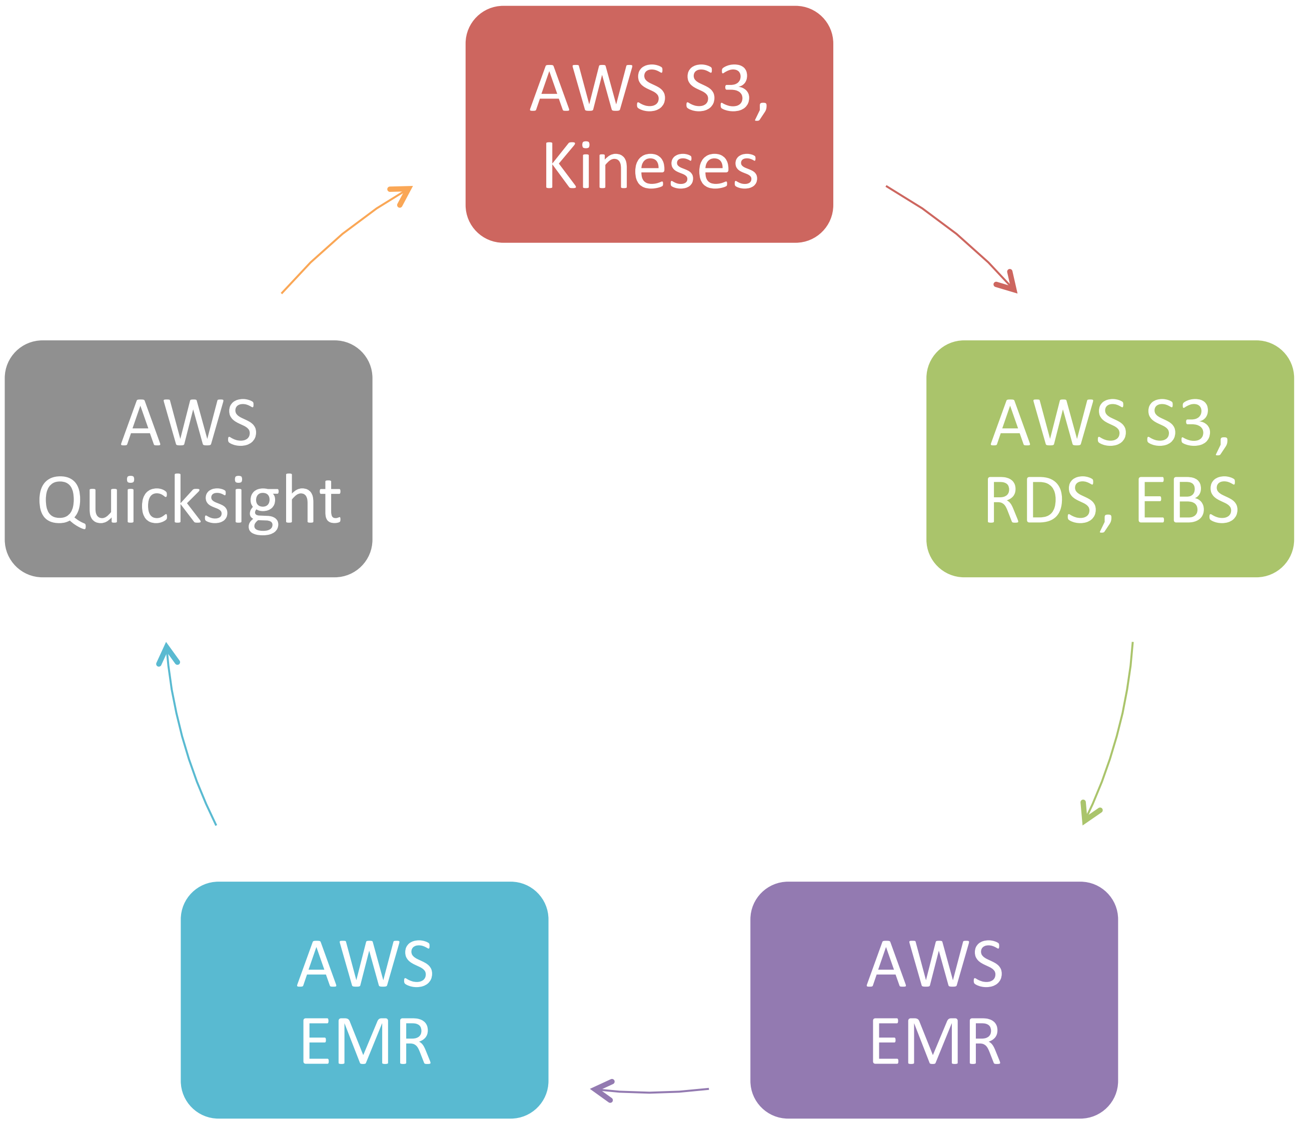

The Life Cycle of Data Analytics

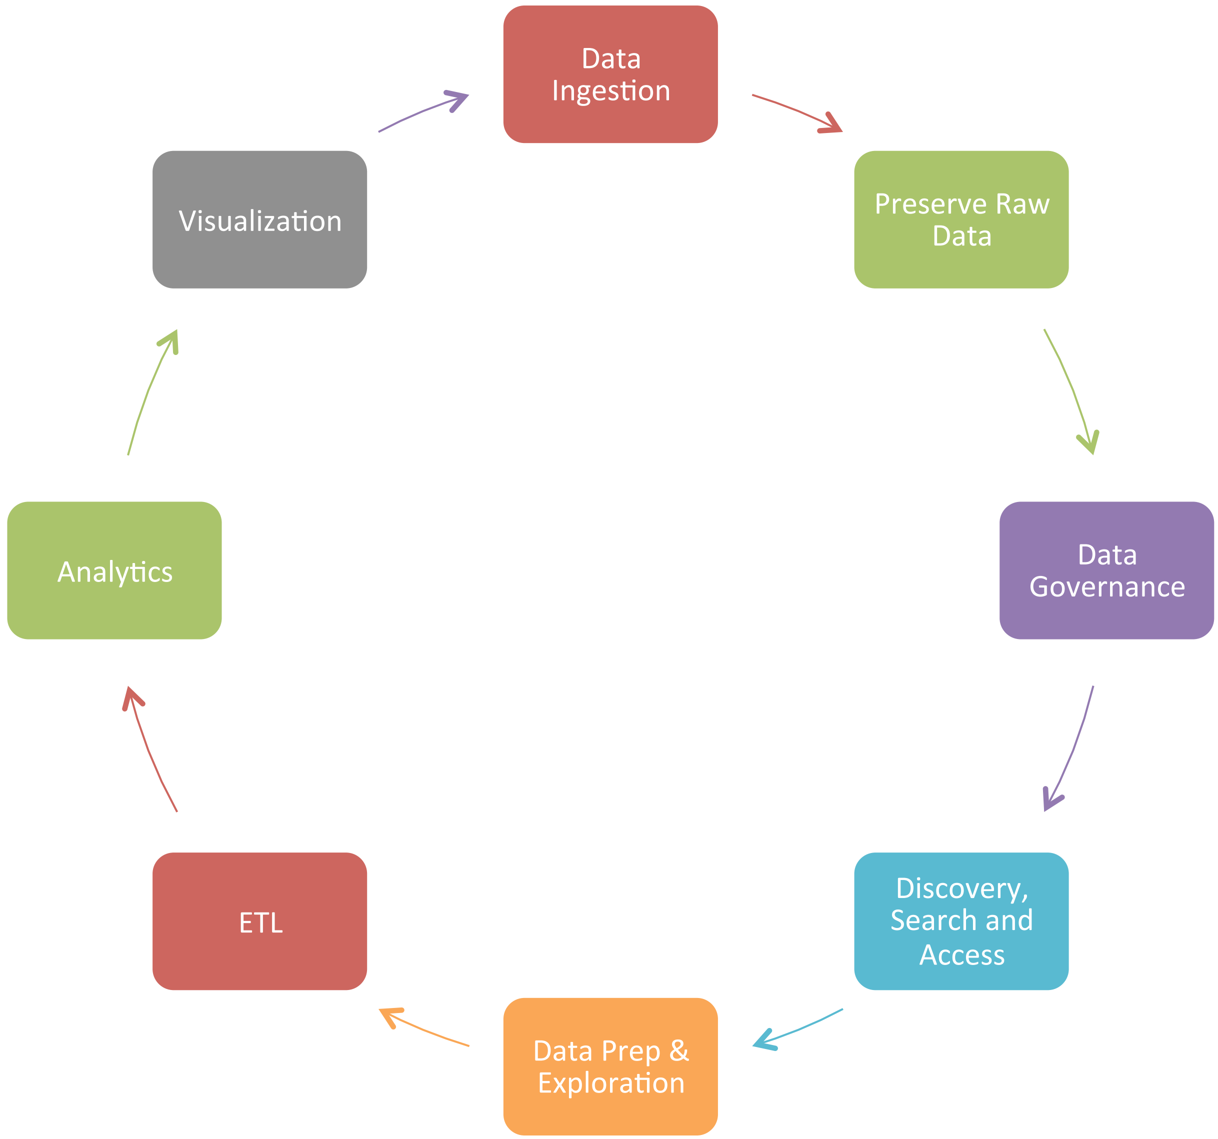

Cloud Use Cases: Analytics

- Cloud services are quickly moving to provide a service for:
  - Software development, testing and deployment
  - Data ingestion, cleaning and analytics

Lambda Architecture

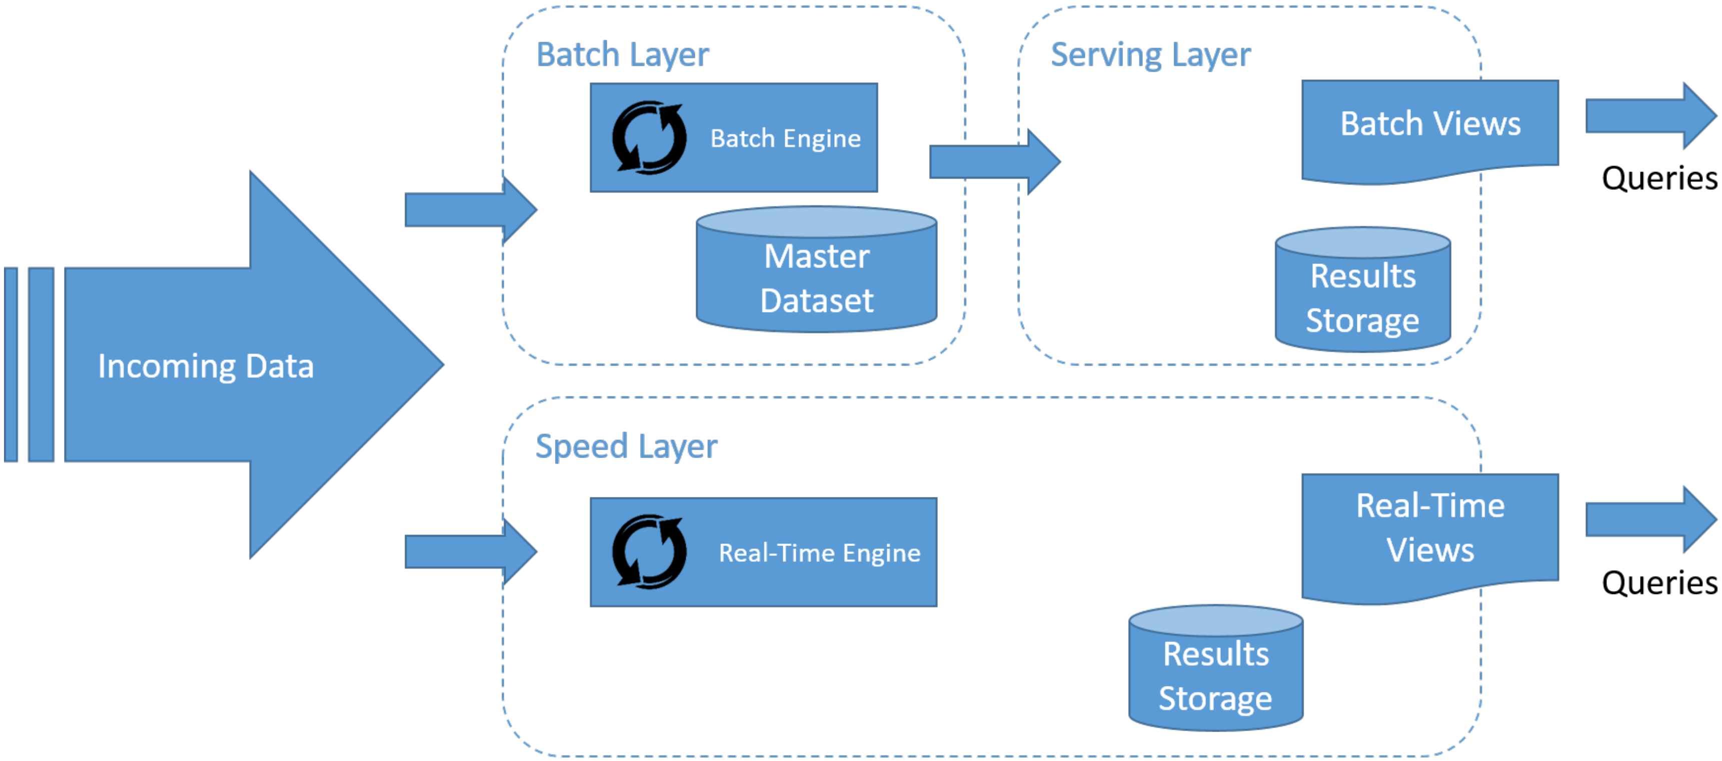

15-719/18-847b
Advanced Cloud Computing

Lecture 02

Cloud Use Cases – Cloud Migration Exercise

January 23, 2017

http://www.cs.cmu.edu/~15719/

A Dept's Research Computing

- Should research departments continue to procure on-premise infrastructure?
- How about migrating the research compute to the cloud?
  - Depends on faculty's research grants
  - Different faculty have different needs
  - Different cloud options to consider
  - Different cost models and SLAs
  - Need to price it based on projected utilization

Baseline Cluster Specs

Cluster = one rack.

- Head node:
  - 1x persistent/reserved instances
  - 730 hours/machine/month (average)
  - 8 cores minimum
  - 30 GB RAM minimum
  - 10T directly accessible storage (Persistent disk/S3)

- Compute nodes
  - 38x persistent/reserved instances
  - 730 hours/machine/month (average)
  - 8 cores
  - 15GB RAM

On Premise

Equipment cost based on a 5-year usable lifecycle:

• Rack	$ 1,295.56
• Power	$ 1,623.91
• Network switch	$ 599.95
• Head node (8 cores, 32GB RAM)	$ 4,158.02
• 10T usable disk for Head Node	$ 1,260.00
• Compute nodes (8 cores, 16G)	
– 38x @$1,882.10 each	$71,519.80
– 13x @$1,882.10 each	$24,467.30
• Total (over 5 years) - 38x dedicated nodes:	$80,457.24
• Total (per month) - 38x dedicated nodes:	$ 1,340.96
• Total (over 5 years) - 13x dedicated nodes:	$33,494.74
• Total (per month) - 13x dedicated nodes:	$ 556.75

Cloud 1

Cost per month:

• Head node (Axlarge)	
– 730 hrs/month x $0.526/hr (reserved)	$383.98
• Storage	
– S3 persistent disk (10T)	$302.60
• Compute (Bxlarge) reserved term (1yr)	
– 38x @730 hrs/month(each) x.419/hr (on-demand)	$11,623.06
– 13x @730 hrs/month(each) x.263/hr (reserved)	$ 2,495.87
– 38x @730 hrs/month(each) x.263/hr (reserved)	$ 7,295.62
– 38x @730 hrs/month(each) x.09/hr (spot)	$ 2,496.60
• Total - based on 100% utilization	
– 38x reserved nodes	$7,982.20
– 13x reserved nodes	$3,182.45

Cloud 2

Cost per month:

• Head node (A-standard-32) 730hrs/month:	$817.60
• Storage - persistent disk (10T)	$410.80
• Compute (B-standard-8)	
– 13x @ 730hrs/month(each)	$759.25
– 38x @ 243hrs/month(each)	$739.73
– 38x @ 730hrs/month(each)	$2,219.35
• Total - based on 100% utilization	
– 38x reserved nodes	$6,395.65
– 13x reserved nodes	$1,987.65

Discussion

- What's missing in this comparison?
  - On Premise:
    - Power, cooling, real estate, personnel, Service/ replacement, safety
  - Cloud:
    - Moving data in/out of cloud, cost of data movement
    - No reuse of compute resources after lifecycle

Private clouds for "enterprises" (established organizations)

15-719/18-847b

Greg Ganger
Garth Gibson
Majd Sakr

Enterprise clusters

- Still common in many places
- Shifting toward being referred to as “private clouds”
  - often not appropriately
- First step: virtualizing machines (VMs)
  - still on per-purpose dedicated clusters
    - Not really “private clouds”
  - reading analyzes characteristics of such clusters
    - low utilization
    - little heterogeneity
    - delays in new deployments

Reminder: promised benefits of cloud computing

- Consolidation
  - Higher server utilization (7-25% -> 70+%)
  - Economies of scale
  - e.g., HP went from 80+ data centers to 6 in early 10s
    - and saved $1B/year… over 60% of total annual expense
- Aggregation
  - One set of experts doing it for many
    - instead of each for themselves
- Rapid deployment
  - Rent/deploy when ready and scale as need
    - rather than specify, buy, deploy, setup, then start

So, many enterprises moving toward private clouds

- Cluster infrastructures shared by many groups/purposes
  - Brings into play the various cloud benefits promised
  - Also brings the challenges we'll be talking about this semester
    - resource scheduling, interference, change management, etc.
- Commonly happening based on org size
  - Biggest (established) companies transitioning fastest
    - Smaller established companies "avoiding change" longer
  - Many new companies (e.g., startups) going public cloud from beginning
- Also a thing: "hybrid cloud"
  - On-premise "private cloud" with overflow demand shifted to public cloud
    - obviously, need to be compatable and have decision policies

Google's infrastructure as a private cloud example

- Huge-scale private cloud
  - Almost everything on huge shared clusters
    - External facing services, dev/test, data analytics, ML, etc.
  - Hardware organized into large-ish clusters (10s to 100s of 1000s of nodes)
    - Scheduled independently, but can shift from one to another
    - Storage accessible within clusters and across (with more overhead)
- Workloads are very heterogeneous and dynamic (and intense)
  - Submitted jobs come in variety of shapes and sizes and constraints
    - inc. "boulders" (many nodes, long-running) and "sand" (smaller, shorter)
  - e.g., 100K+ scheduler decisions/hour in 12.5K-node cluster
- Resources heterogeneous and dynamic too
  - Different machine generations, accelerators, locations, partially-failed, etc.

Next day plan

- Prepare to discuss OpenStack and building our own cloud