DESIGN, AUTOMATION & TEST IN EUROPE
27 – 31 March, 2017 · STCC · Lausanne · Switzerland
The European Event for Electronic System Design & Test

Server Benchmarking and Design with CloudSuite 3.0

Javier Picorel

PARALLEL SYSTEMS ARCHITECTURE LAB

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

A Brief History of IT

Mobile Era

Consumer Era

1970s- | 1980s | 1990s | Today+

Mainframes

PC Era

WWW Era

- From computing-centric to data-centric
- Consumer Era: Internet-of-Things in the Cloud

Data is Shaping Future of IT

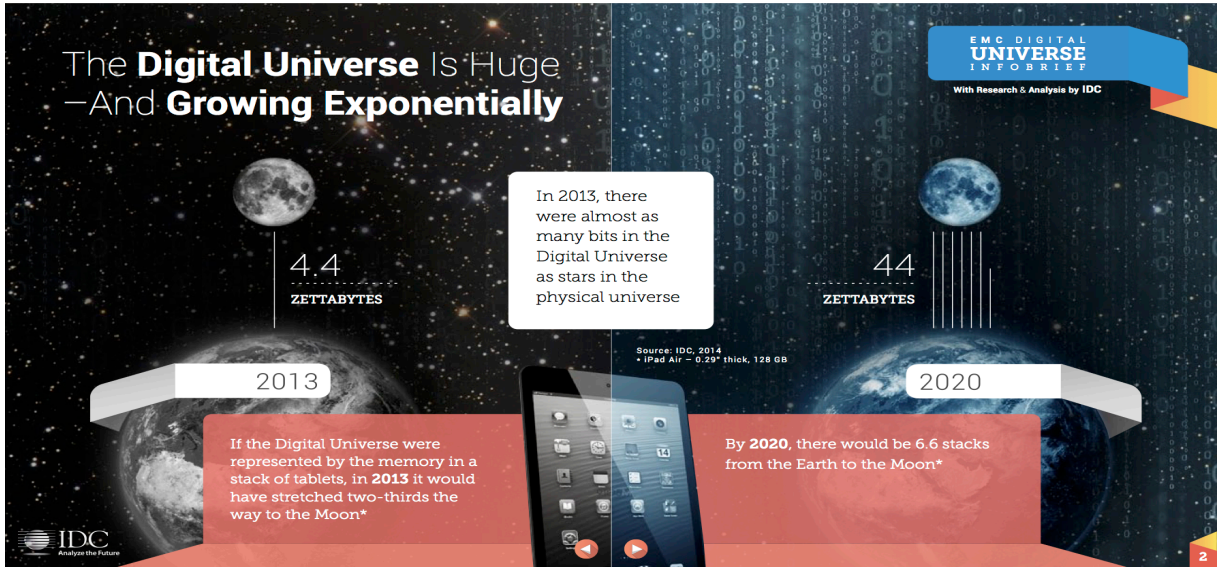

- Data growth = 100x in ten years [IDC 2015]
- Monetizing data for commerce, health, science, services, ….
- Big Data is shaping IT & pretty much whatever we do!

Data Shaping All Science & Technology

Science (traditionally HPC) entering 4th paradigm

- Analytics using IT on
  - Instrument data
  - Simulation data
  - Sensor data
  - Human data

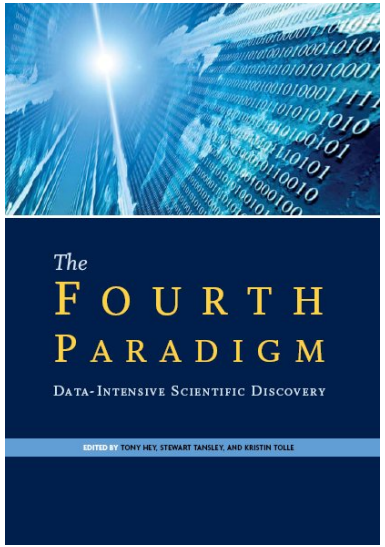

Complements empirical science, theory & simulation

HPC & data-centric cloud services are converging

Modern HPC in the Datacenter

- Increasing popularity of analytics workloads
  - Closely related to traditional HPC workloads (e.g., graphs)
- Service providers don't acquire supercomputers
  - All workloads share the same datacenter
  - Cost hard to sustain (e.g., IBM discontinued BlueGene)
- HPC is taking a turn towards datacenters
  - Datacenter provides higher availability, lower queue times, flexibility
  - E.g., Amazon provides HPC instances

Datacenters are the heart of both cloud services and science

Datacenters Growing Rapidly

Source: James Hamilton, 2012

Daily growth in 2012 = First five years of business!

How are we Designing Cloud Systems?

- Spoiler alert: We are doing it wrong!
- Modern servers based on desktop processors
- Server design guided by unrepresentative benchmarks

Design needs to be driven by cloud-representative benchmarks

Traditional Benchmarks

- SPEC, PARSEC, TPC-C, SPLASH, …
  - Single machine metrics (e.g., SPEC score)
  - Metrics measure performance of a single component (usually CPU)
  - Vastly different application footprints

- None of these run on a datacenter

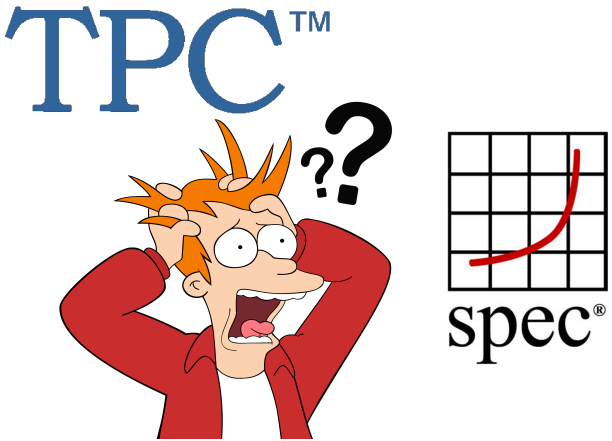

Traditional benchmarks not suitable for cloud evaluation and research

What prior benchmarks miss?

- No notion of end-to-end performance metrics
  - Traditional metrics do not reflect user experience
- Cloud services have extensive instruction footprint
  - Multi-megabyte instruction working sets
  - Overall performance highly dependent on processor's frontend
- Cloud services deal with big data
  - Datasets do not fit in on-chip caches

What those benchmarks are modeling, does not apply here

Cloud Service Requirements

- Throughput:
  Owners want computing capabilities for their money

- Latency (online services):
  Users abandon services if response time is high
  - Amazon: 100ms of latency can cause 1% of sales loss
  - Google Search: 500ms of latency dropped traffic by 20%
  - YouTube: Users start abandoning video after 2 seconds of wait

CloudSuite Benchmark Suite

- CloudSuite's goal:
  Assess performance of cloud services on modern hardware
  - Make the case for cloud service representativeness
  - Identify improvement opportunities for server hardware

- End-to-end performance metrics
  - Hard problem; still online service dependent!

Cloud Benchmarking with CloudSuite 3.0

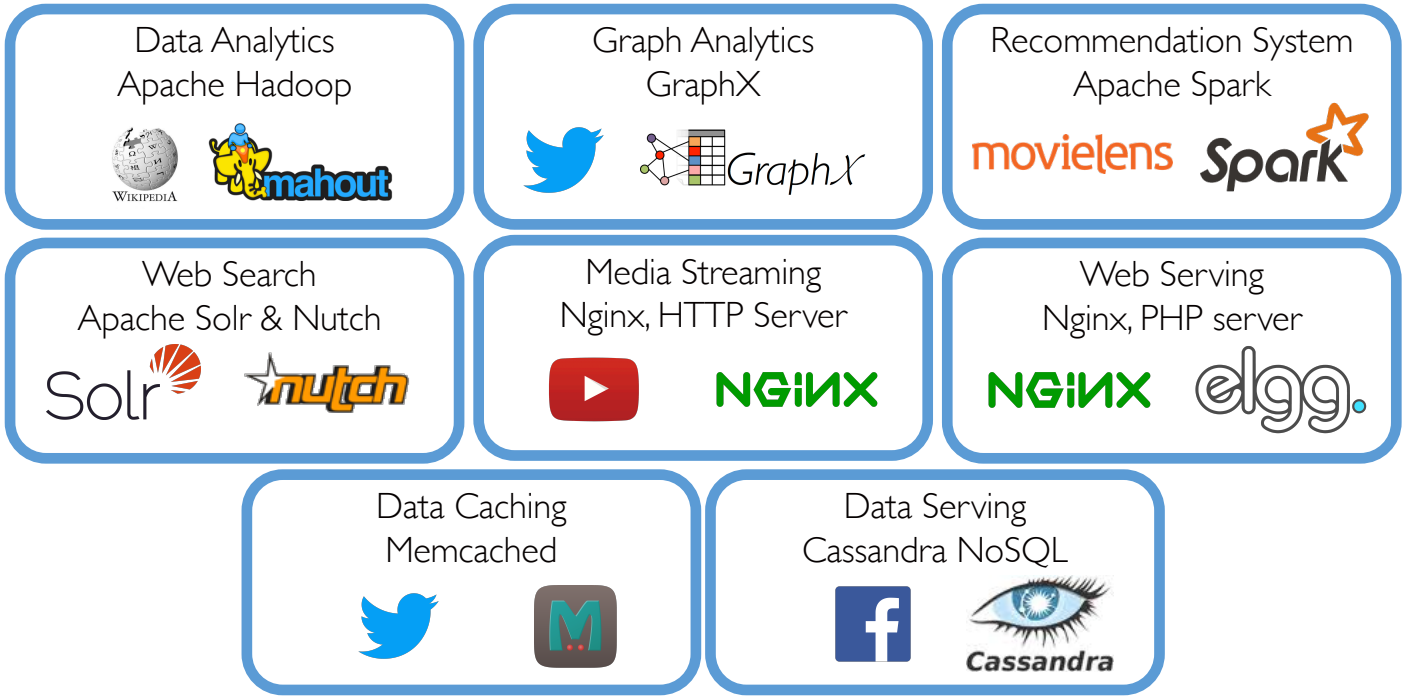

CloudSuite

Available in cloudsuite.ch

Brief History of CloudSuite

- Clearing the Clouds [Ferdman et al., ASPLOS'12] (CloudSuite 1.0)
  - Fundamental mismatch of cloud workloads and modern servers
  - Sever silicon real-estate misuse in current systems
- CloudSuite 2.0 – two additional benchmarks
  - Graph Analytics, Data Caching
- Insights derived from CloudSuite impacted industry
  - E.g., Cavium ThunderX
- CloudSuite 3.0 

  - Integration with Docker & Google's PerfKit Benchmarker
  - Additional benchmark: Recommendation system

Clearing the Clouds in a Nutshell [ASPLOS 2012]

Workload/Server Mismatch

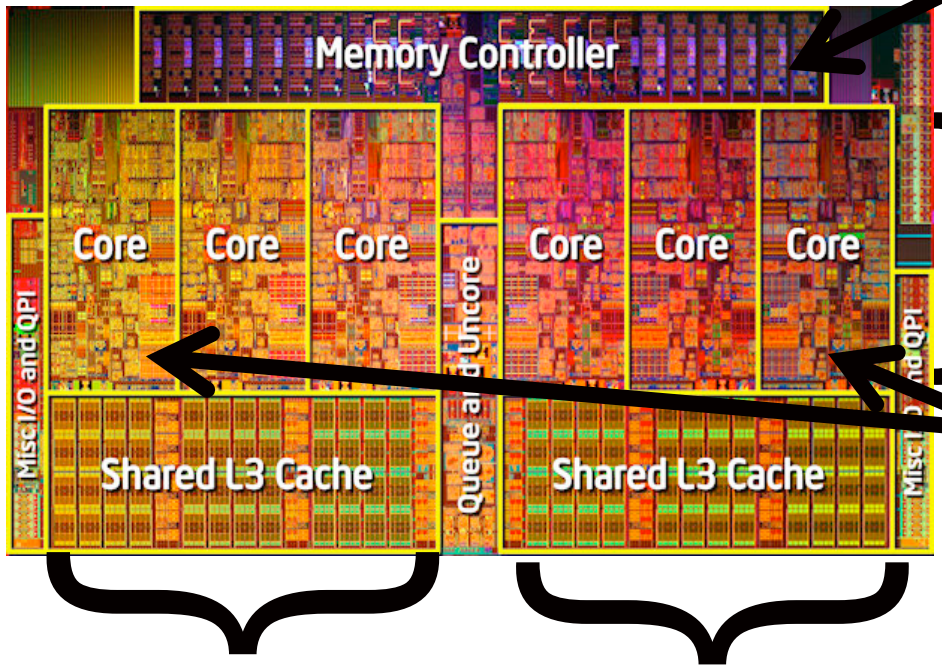

B/W unused!

Cores too fat!

Too few cores!

20 MB (60%) waste of space (no reuse)!

Insights Impacted Industry

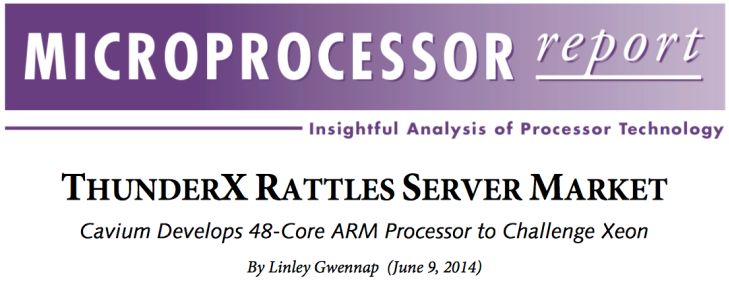

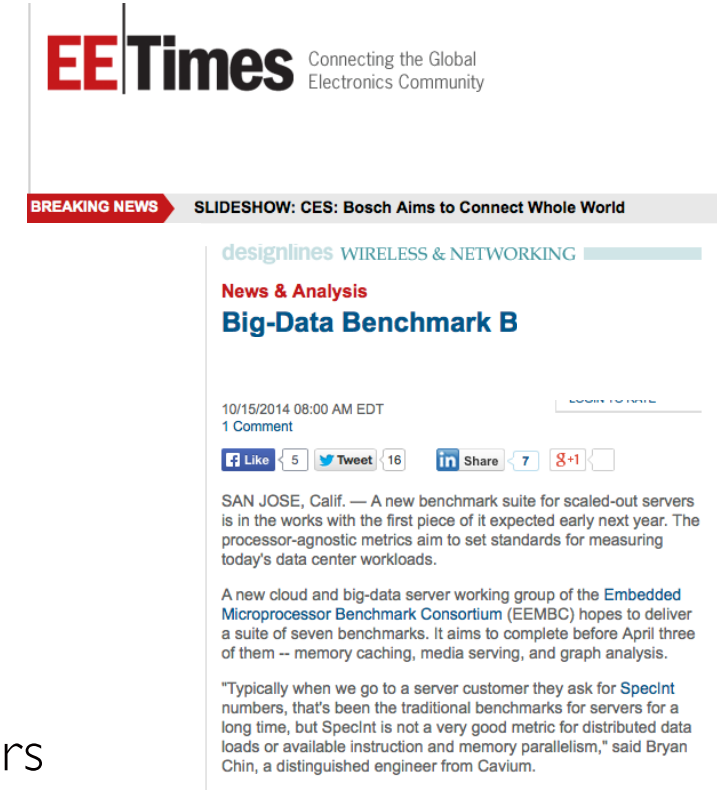

48-core 64-bit ARM SoC

[Influenced industry products]:

- Designed to serve data
- Specialized chip design for servers
- 10x better efficiency than Xeon

Google PerfKit Benchmarker

- Goal: Standardize Cloud performance evaluation
- A tool to compare cloud service providers
- Consortium of industry and academics

- Automates benchmarks including creating databases, disks, VMs, …
  - 26+ benchmarks
  - CloudSuite 3.0 benchmarks included

- Shared publicly on GitHub
  - http://www.github.com/GoogleCloudPlatform/PerfKitBenchmarker

Perfkit's Workflow

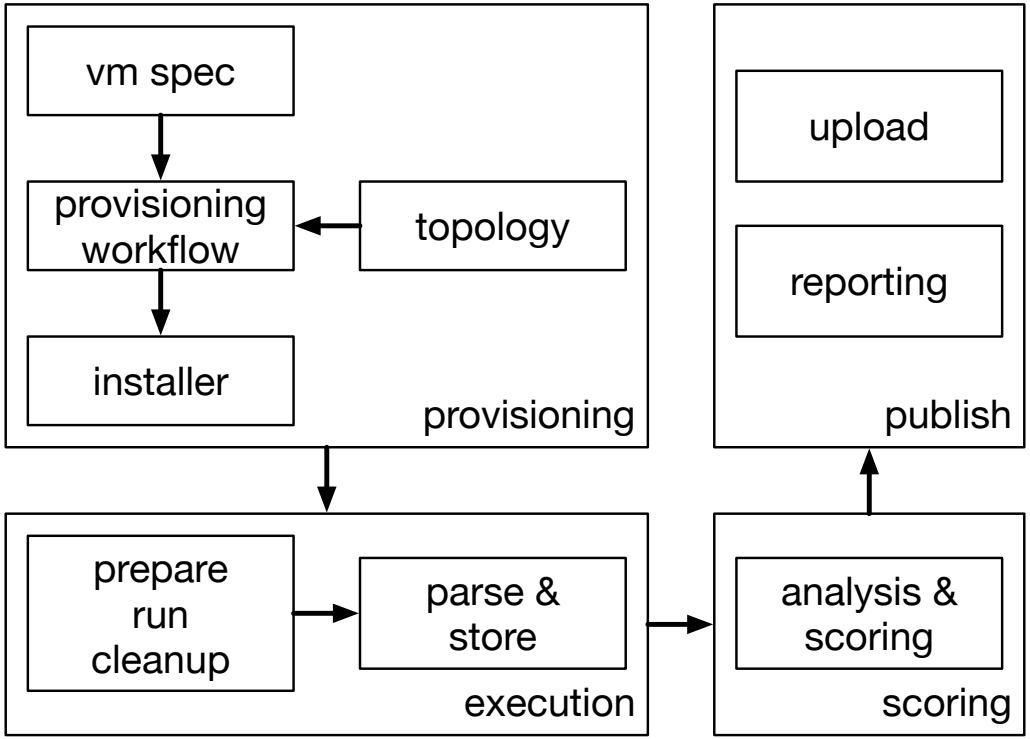

Perfkit automates the deployment and benchmarking processes

What's new in CloudSuite 3.0

- A couple of different workloads
  - New benchmark: Recommendation system
  - New software stacks: Graph Analytics, Media Streaming, Web Search
- Updated software packages of all workloads
- Docker containers → ease of deployment
  - This is huge! (literally)

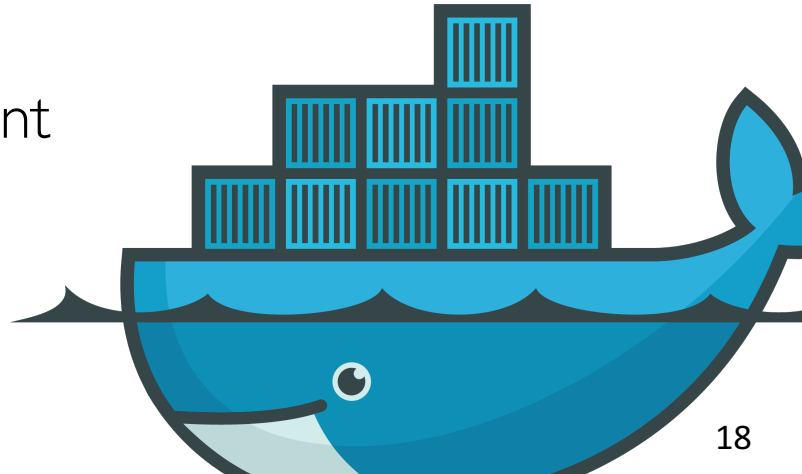

Target Audience

- System designers
  - Assess & compare systems' performance of cloud workloads
- Computer architects
  - Derive insights for future server design
- HPC community
  - Datacenter & HPC applications converging

Key Cloud Service Characteristics

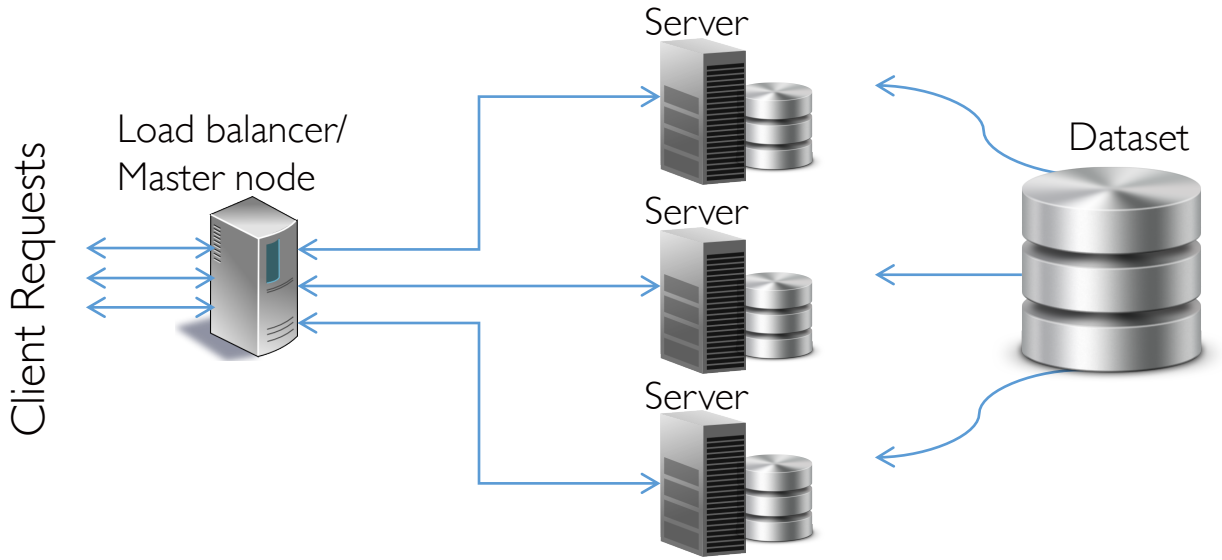

- Serve independent requests/tasks
- Operate on huge dataset split into shards
- Communicate infrequently or in bulk
- Strict real-time constraints (for online services)

CloudSuite 3.0 Benchmarks

- Offline benchmarks
  - Data Analytics
  - Recommendation system
  - Graph Analytics

- Online benchmarks
  - Data Caching
  - Data Serving
  - Media Streaming
  - Web Search
  - Web Serving

Offline Benchmarks

- Operate on large datasets
- Usually a machine learning algorithm over large datasets
- Performance metric:
  - Completion time (for a given input size)
- No real-time constraints

Data Analytics

- Massive amounts of human-generated data (Big Data)
- Extract useful information from data
  - Predict user preferences, opinions, behavior
  - Benefit from information (e.g., business, security)
- Several examples
  - Book recommendation (Amazon)
  - Spyware detection (Facebook)

Data Analytics Benchmark

- **Application:** Text classification
  - Sentiment analysis
  - Spam Identification

- **Software:** Mahout (Apache)
  - Popular MapReduce machine learning library

- **Dataset:** Wikipedia English page articles

Data Analytics Benchmark

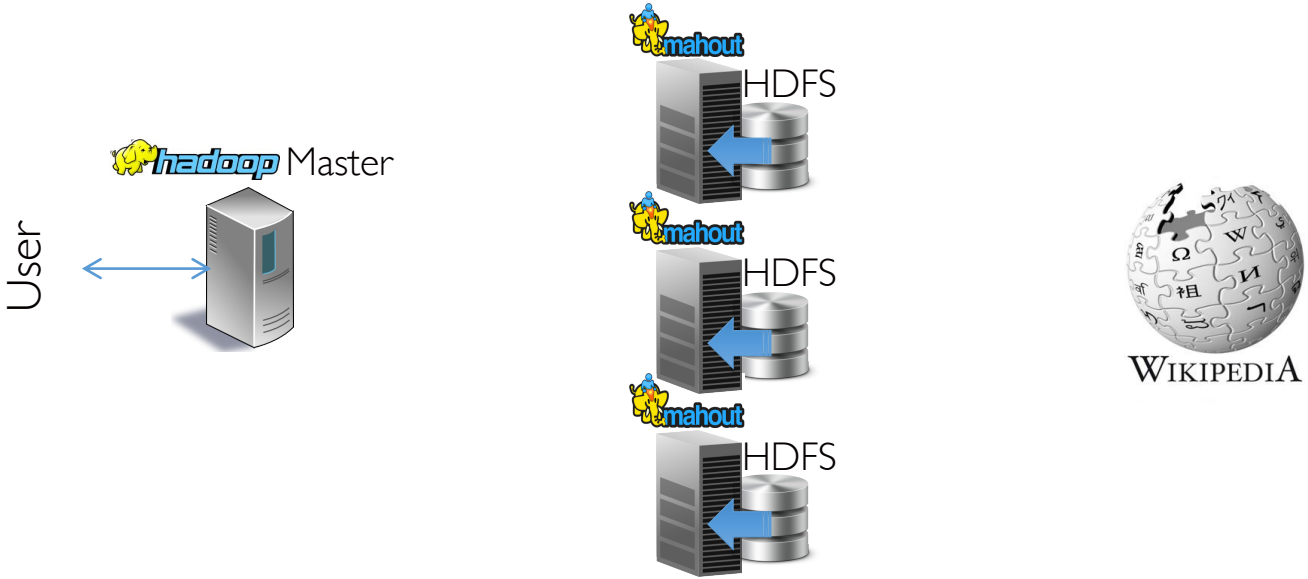

- Build a model from a Wikipedia training input
- Master sends Wikipedia documents for classification
- Slaves classify documents locally using model and send results to master
- Performance metric: completion time

CloudSuite 3.0 Benchmarks

- Offline benchmarks
  - Data Analytics
  - Recommendation system
  - Graph Analytics

- Online benchmarks
  - Data Caching
  - Data Serving
  - Media Streaming
  - Web Search
  - Web Serving

Recommendation System

- Extract useful information from user data
  - Predict user preferences, rates

- Several examples
  - Movie recommendation (Netflix)
  - Item recommendation (Amazon)
  - Song recommendation (Spotify)
  - Recommending new friends, groups, … (Social networks)

NETFLIX

amazon

Recommendation System Benchmark

- **Application:** Collaborative filtering
  - Recommendation systems

- **Software:** Apache MLlib
  - Popular Apache Spark machine learning library

- **Dataset:** Movielens video database

Recommendation System Benchmark

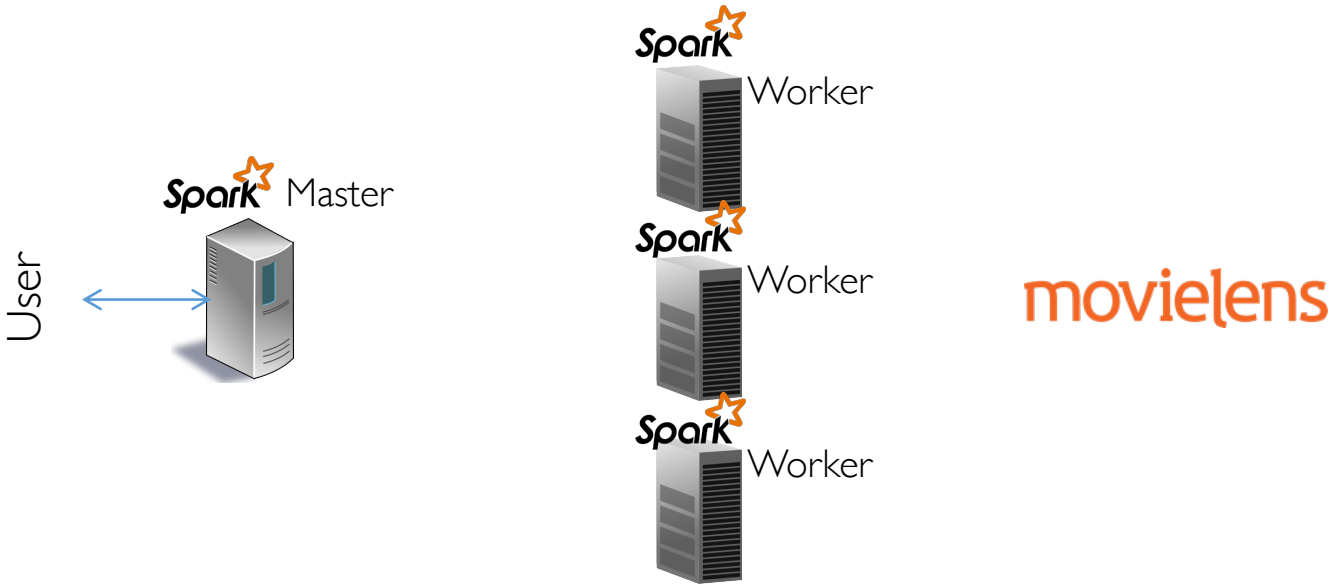

- Build a recommendation model with ALS matrix factorization
- Master partitions rating matrix, user & item vectors; sends them to workers
- Workers perform local matrix factorization and send results to master
- Performance metric: completion time

CloudSuite 3.0 Benchmarks

- Offline benchmarks
  - Data Analytics
  - Recommendation system
  - Graph Analytics

- Online benchmarks
  - Data Caching
  - Data Serving
  - Media Streaming
  - Web Search
  - Web Serving

Graph Analytics

- Parallel distributed graph processing
- Data mining on graphs
- Graph examples
  - Social networks (e.g., Facebook, Twitter)
  - Web graph

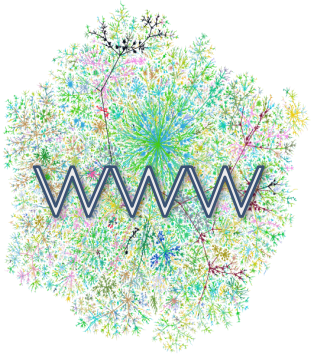

Graph Analytics Benchmark

- Application: PageRank
  - Measures influence of Twitter users
  - How much attention followers pay to a user

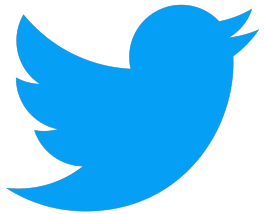

- Software: Apache GraphX
  - Parallel framework for graph processing

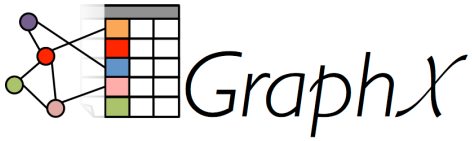

- Dataset
  - Twitter user graph

Graph Analytics Benchmark

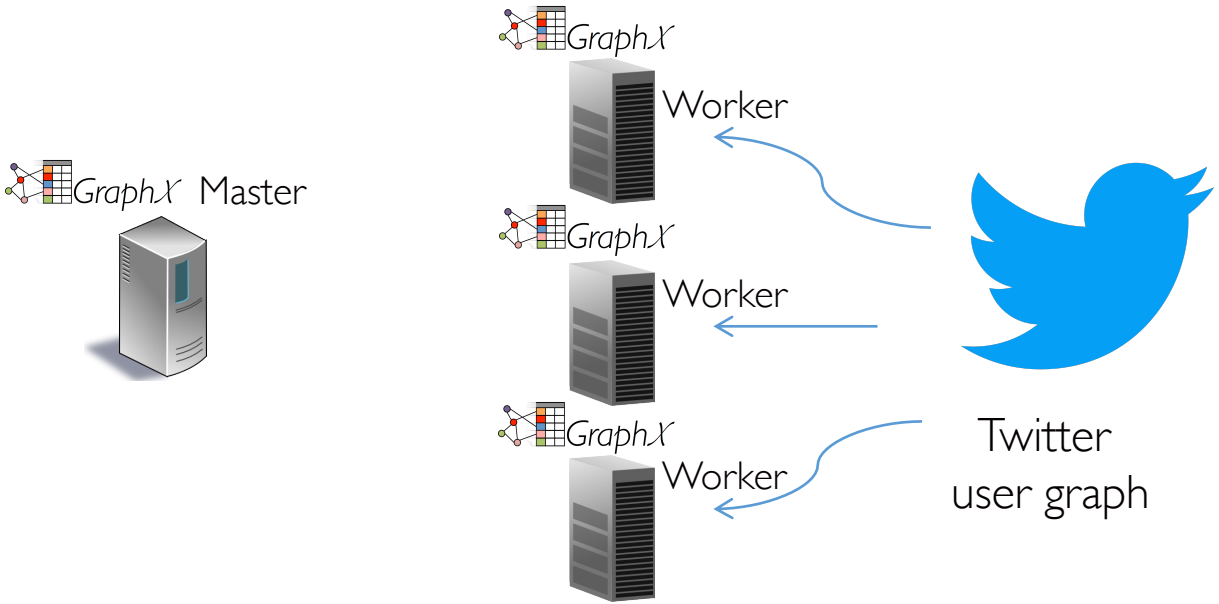

- Distributes the graph across nodes
- Iterative computation: Always with adjacent vertices
- Communication across machines for adjacent vertices
- Output: influence of each user in the graph
- Performance metric: completion time

CloudSuite 3.0 Benchmarks

- Offline benchmarks
  - Data Analytics
  - Recommendation system
  - Graph Analytics
- Online benchmarks
  - Data Caching
  - Data Serving
  - Media Streaming
  - Web Search
  - Web Serving

Online Benchmarks

- Operate on large datasets

- Throughput is important, but also need high service quality
  - Tail latency of requests is critical for service quality
  - Goal: Maximize throughput *under QoS target*

- Performance metrics:
  - Throughput (metric is benchmark-specific)
  - Delivered QoS (in terms of N-th percentile latency)

Data Caching

- Web apps are latency-sensitive
- Fetching data from disk is slow
- Caching data in memory for fast data access
  - General-purpose, in-memory key-value store
  - Caches data for other apps, another tier before back-end

Data Caching Benchmark

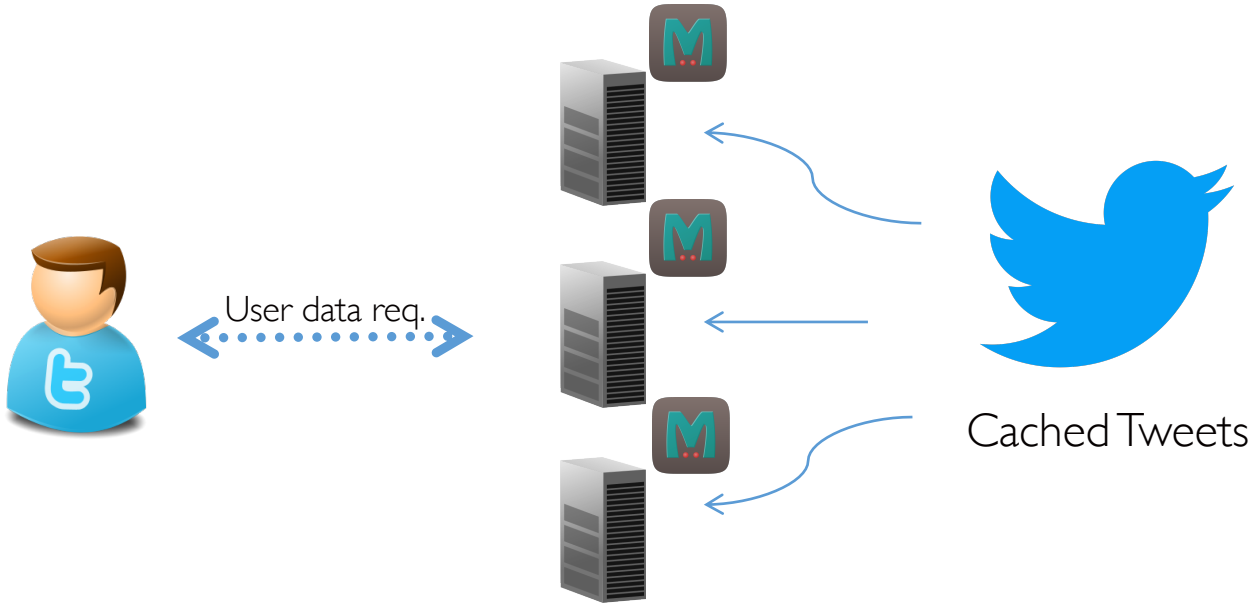

- Driver emulates Twitter users
- Memcached software to cache data in memory
- If data not found in cache, returns a miss reply
- Performance metrics: # requests/second, N-th pct latency

CloudSuite 3.0 Benchmarks

- Offline benchmarks
  - Data Analytics
  - Recommendation system
  - Graph Analytics
- Online benchmarks
  - Data Caching
  - Data Serving
  - Media Streaming
  - Web Search
  - Web Serving

Data Serving

- Global-scale online services rely on NoSQL datastores
  - Inherently scalable
  - Suitable for unpredictable schema changes
- Scale out to meet service requirements
  - Accommodate fast data generation rate

Data Serving Operation

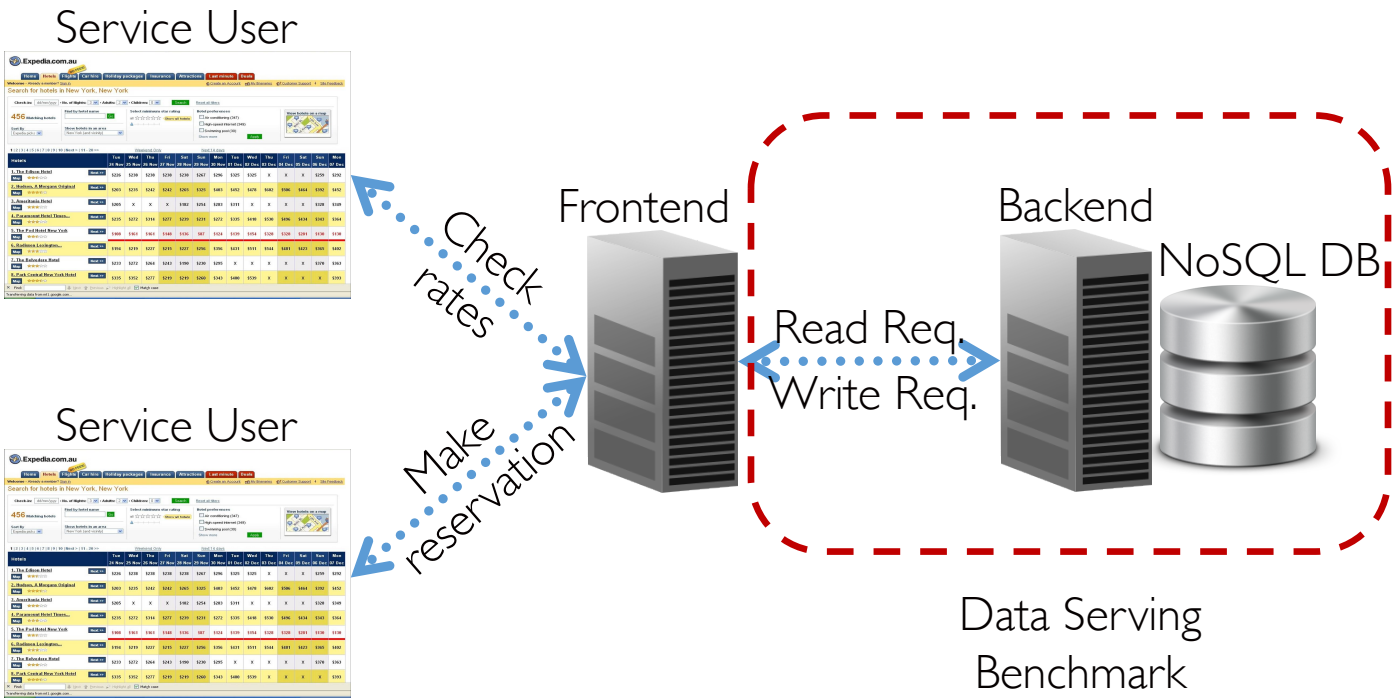

Data Serving Benchmark

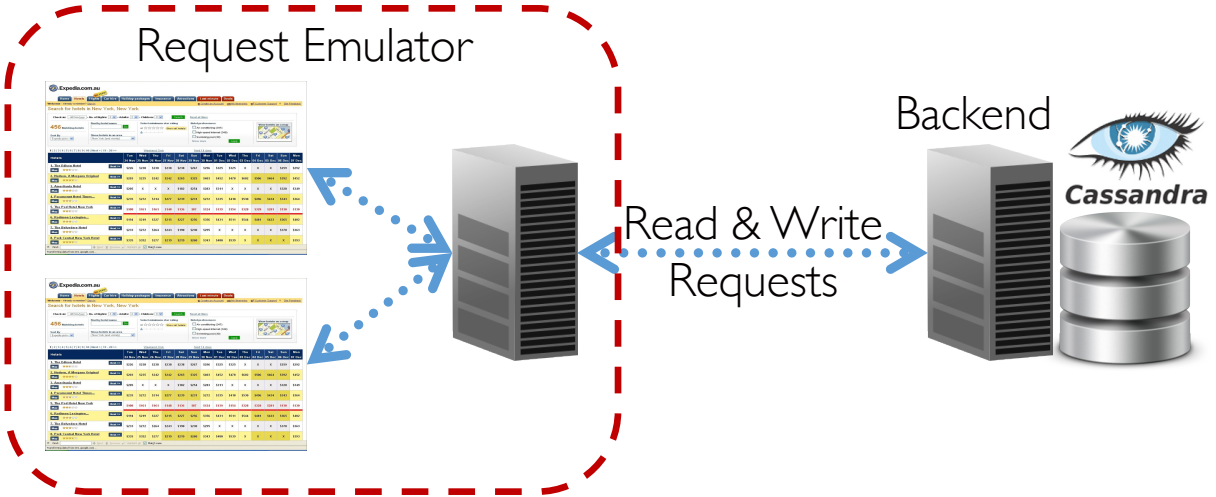

- Yahoo! Cloud Serving Benchmark (YCSB) driver
  - Predefined mixes of read/write operations
  - Popularity of access distributions (e.g., zipfian)
  - Interface to popular datastores (e.g., Cassandra, HBase)

Data Serving Benchmark

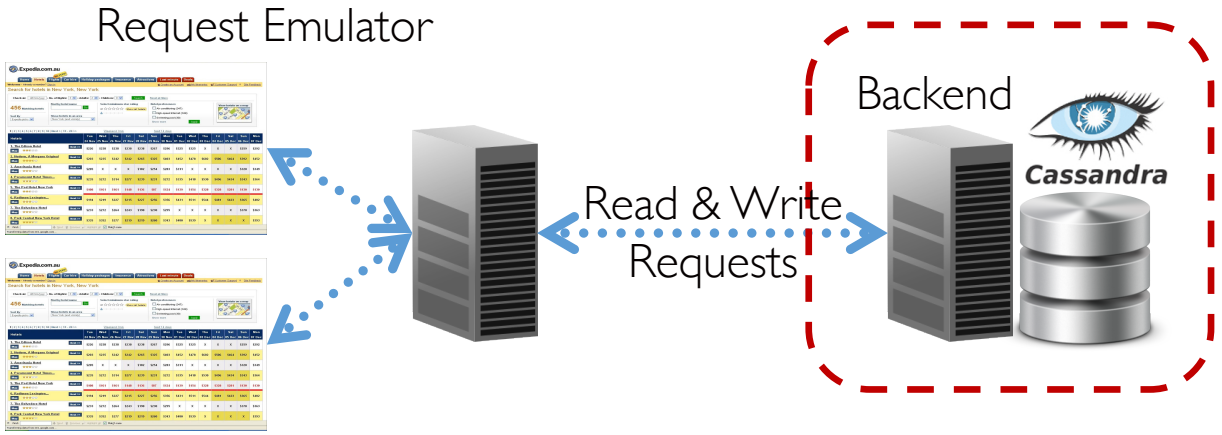

- Cassandra datastore
  - Popular NoSQL: many use cases (e.g., Expedia, eBay, Netflix)
- Driver generates dataset
  - Defines number & size of fields
  - Populates datastore
- Performance metrics: R/W ops/s, N-th pct latency

CloudSuite 3.0 Benchmarks

- Offline benchmarks
  - Data Analytics
  - Recommendation system
  - Graph Analytics
- Online benchmarks
  - Data Caching
  - Data Serving
  - Media Streaming
  - Web Search
  - Web Serving

Media Streaming

- Media streaming expected to dominate internet traffic

- Increasing popularity of media streaming services
  - Video sharing sites, movie streaming services, etc.

Media Streaming Operation

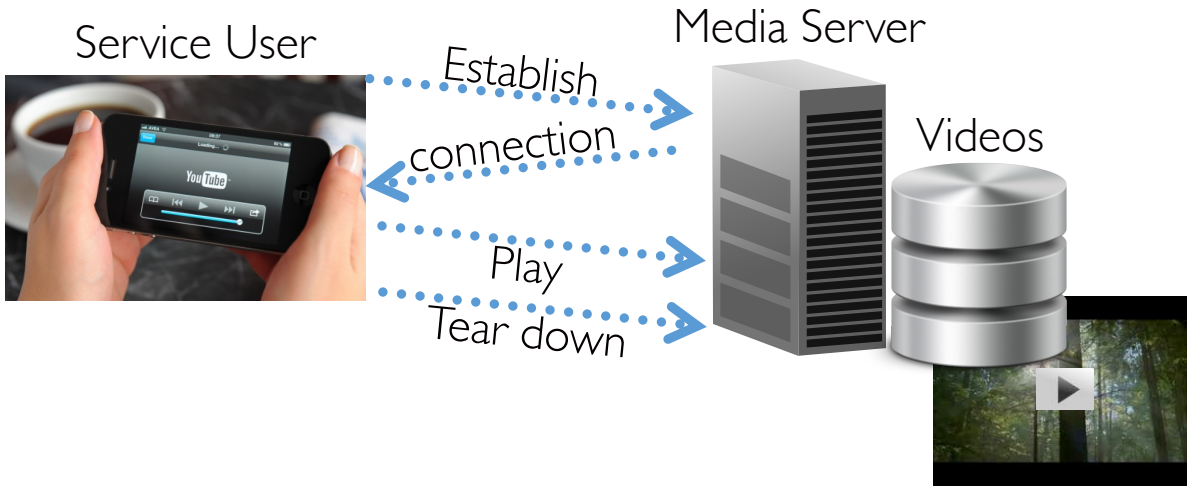

Media Streaming Benchmark

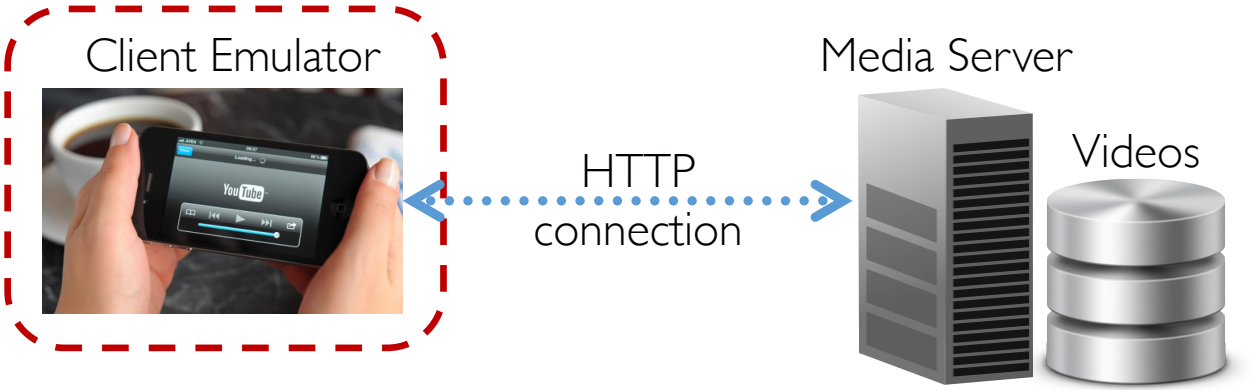

- Implements HTTP communication
- Uses the videoperf client, based on the httperf traffic generator
- Allows a flexible mix of requests
  - Different video lengths and qualities

Media Streaming Benchmark

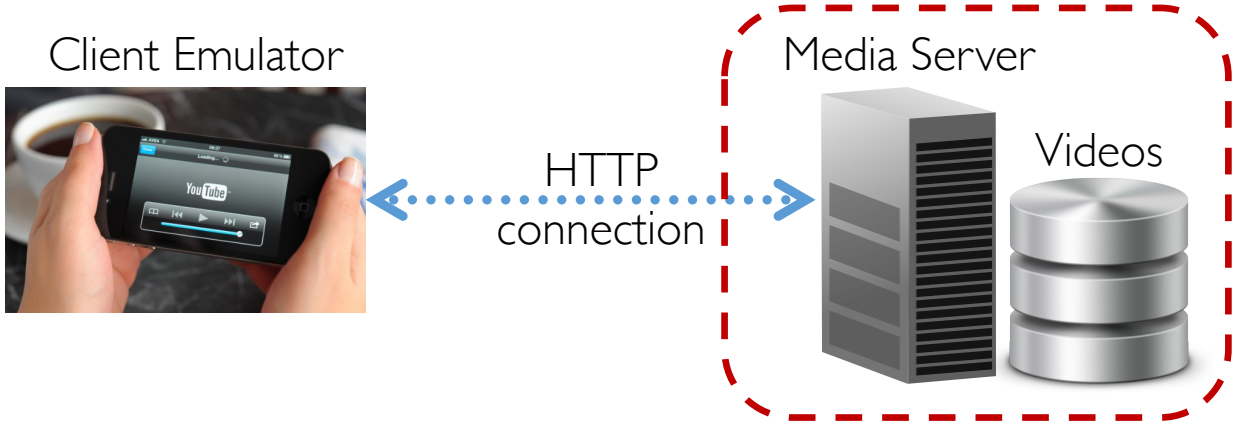

- Server required to support HTTP
  - Nginx server
- Dataset consists of a mix of pre-encoded videos
  - Four video qualities of different durations (240p, 360p, 480p, 720p)
  - Exponential popularity distribution
- Performance metrics: streaming bandwidth (Kbps), avg. reply delay

CloudSuite 3.0 Benchmarks

- Offline benchmarks
  - Data Analytics
  - Recommendation system
  - Graph Analytics
- Online benchmarks
  - Data Caching
  - Data Serving
  - Media Streaming
  - Web Search
  - Web Serving

Web Search

- Most popular online service
  - Numerous search engines deployed by industry

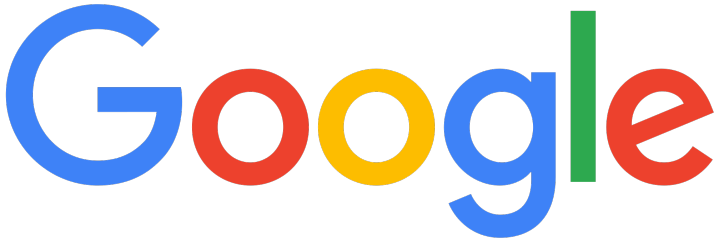

Web Search Operation

Search User

Frontend

Query Term	Document
...	
Benchmark	1, 10, 17, ...
CloudSuite	3, 45, ...
Datacenter	9, 11, 14, 45, ...
EPFL	17, 10, 15, ...
PerfKit	3, 4,18
...	

Index Serving Node (ISN)

Query Term	Document
...	
Benchmark	1, 5, 7, ...
CloudSuite	5, 2, ...
Datacenter	7, 10, 17, 20, ...
EPFL	2, 4, 6, 8, 23, ...
PerfKit	3, 5, 20, 33, 34, 55, ...
...	

Inverted Index

Query Term	Document
...	
Benchmark	1, 6, 19, ...
CloudSuite	5, 40, ...
Datacenter	6, 10, 13, 20, ...
EPFL	5, 10, 23, ...
PerfKit	3, 6, 10, 20, ...
...	

Web Search Operation

Search User

Query
= "EPFL"

Frontend

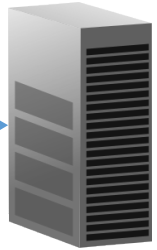

Query Term	Document
...	
Benchmark	1, 10, 17, ...
CloudSuite	3, 45, ...
Datacenter	9, 11, 14, 45, ...
EPFL	17, 10, 15, ...
PerfKit	3, 4,18
...	

Index Serving Node (ISN)

Query Term	Document
...	
Benchmark	1, 5, 7, ...
CloudSuite	5, 2, ...
Datacenter	7, 10, 17, 20, ...
EPFL	2, 4, 6, 8, 23, ...
PerfKit	3, 5, 20, 33, 34, 55, ...
...	

Inverted Index

Query Term	Document
...	
Benchmark	1, 6, 19, ...
CloudSuite	5, 40, ...
Datacenter	6, 10, 13, 20, ...
EPFL	5, 10, 23, ...
PerfKit	3, 6, 10, 20, ...
...	

Web Search Operation

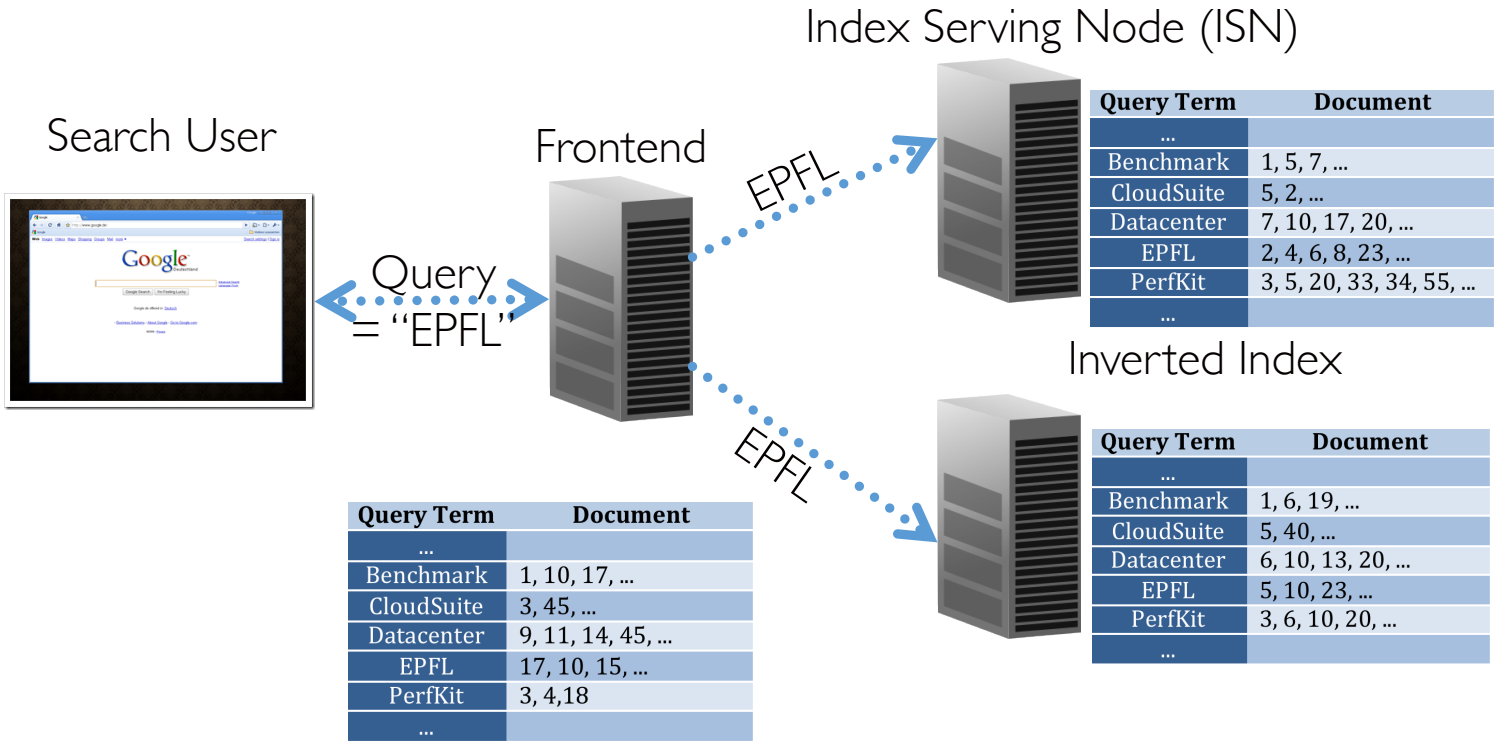

Query Term	Document
...	
Benchmark	1, 5, 7, ...
CloudSuite	5, 2, ...
Datacenter	7, 10, 17, 20, ...
EPFL	2, 4, 6, 8, 23, ...
PerfKit	3, 5, 20, 33, 34, 55, ...
...	

Query Term	Document
...	
Benchmark	1, 6, 19, ...
CloudSuite	5, 40, ...
Datacenter	6, 10, 13, 20, ...
EPFL	5, 10, 23, ...
PerfKit	3, 6, 10, 20, ...
...	

Query Term	Document
...	
Benchmark	1, 10, 17, ...
CloudSuite	3, 45, ...
Datacenter	9, 11, 14, 45, ...
EPFL	17, 10, 15, ...
PerfKit	3, 4,18
...	

Web Search Operation

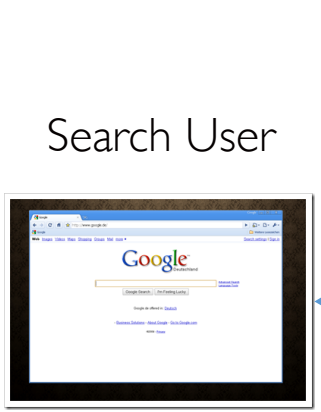

Query = "EPFL"

Frontend

EPFL

Doc. ids

Doc. ids

EPFL

Index Serving Node (ISN)

Query Term	Document
...	
Benchmark	1, 5, 7, ...
CloudSuite	5, 2, ...
Datacenter	7, 10, 17, 20, ...
EPFL	2, 4, 6, 8, 23, ...
PerfKit	3, 5, 20, 33, 34, 55, ...
...	

Inverted Index

Query Term	Document
...	
Benchmark	1, 6, 19, ...
CloudSuite	5, 40, ...
Datacenter	6, 10, 13, 20, ...
EPFL	5, 10, 23, ...
PerfKit	3, 6, 10, 20, ...
...	

Query Term	Document
...	
Benchmark	1, 10, 17, ...
CloudSuite	3, 45, ...
Datacenter	9, 11, 14, 45, ...
EPFL	17, 10, 15, ...
PerfKit	3, 4,18
...	

Web Search Operation

Search User

Query
= "EPFL"

Frontend

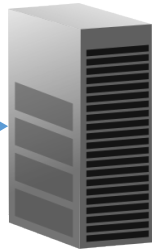

Query Term	Document
...	
Benchmark	1, 10, 17, ...
CloudSuite	3, 45, ...
Datacenter	9, 11, 14, 45, ...
EPFL	17, 10, 15, ...
PerfKit	3, 4,18
...	

Index Serving Node (ISN)

Query Term	Document
...	
Benchmark	1, 5, 7, ...
CloudSuite	5, 2, ...
Datacenter	7, 10, 17, 20, ...
EPFL	2, 4, 6, 8, 23, ...
PerfKit	3, 5, 20, 33, 34, 55, ...
...	

Inverted Index

Query Term	Document
...	
Benchmark	1, 6, 19, ...
CloudSuite	5, 40, ...
Datacenter	6, 10, 13, 20, ...
EPFL	5, 10, 23, ...
PerfKit	3, 6, 10, 20, ...
...	

Web Search Operation

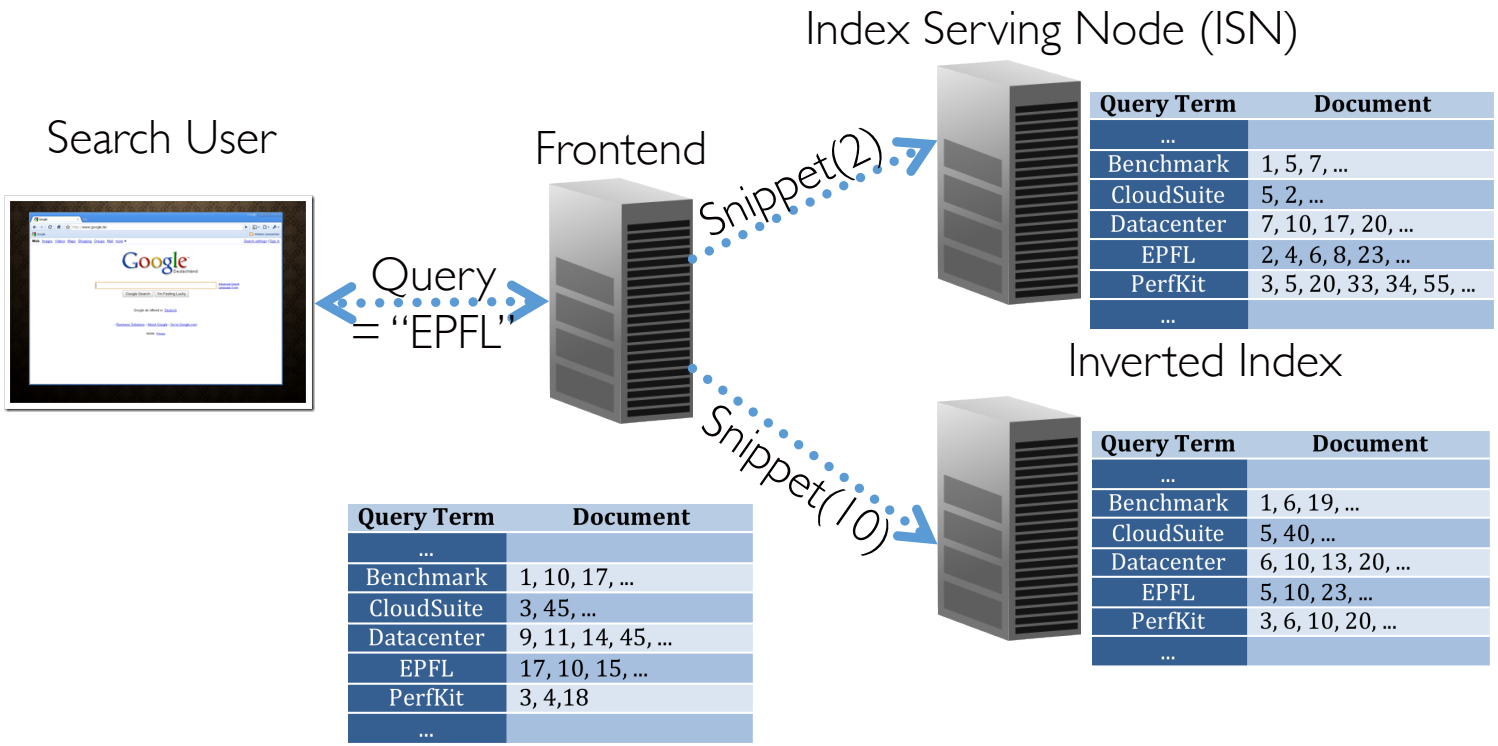

Web Search Operation

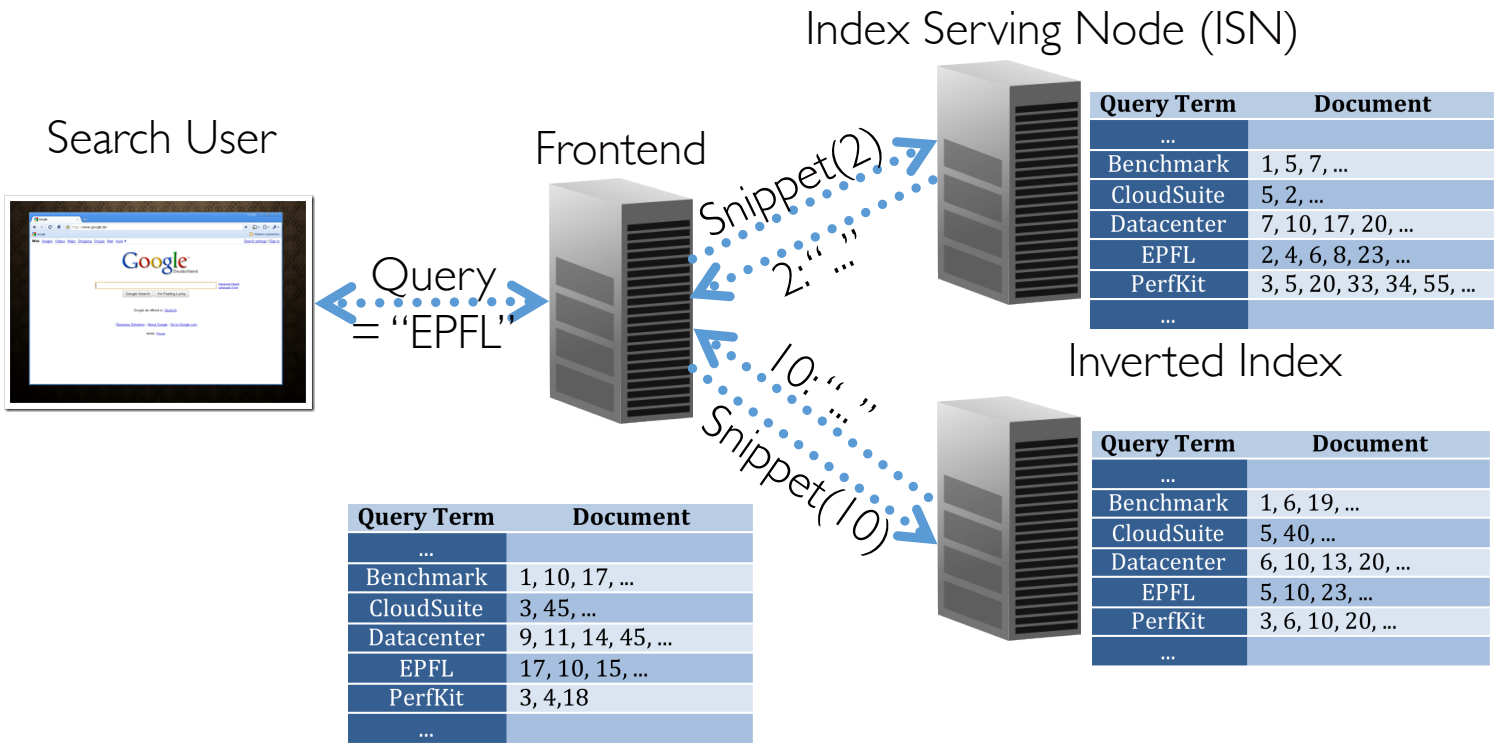

Web Search Operation

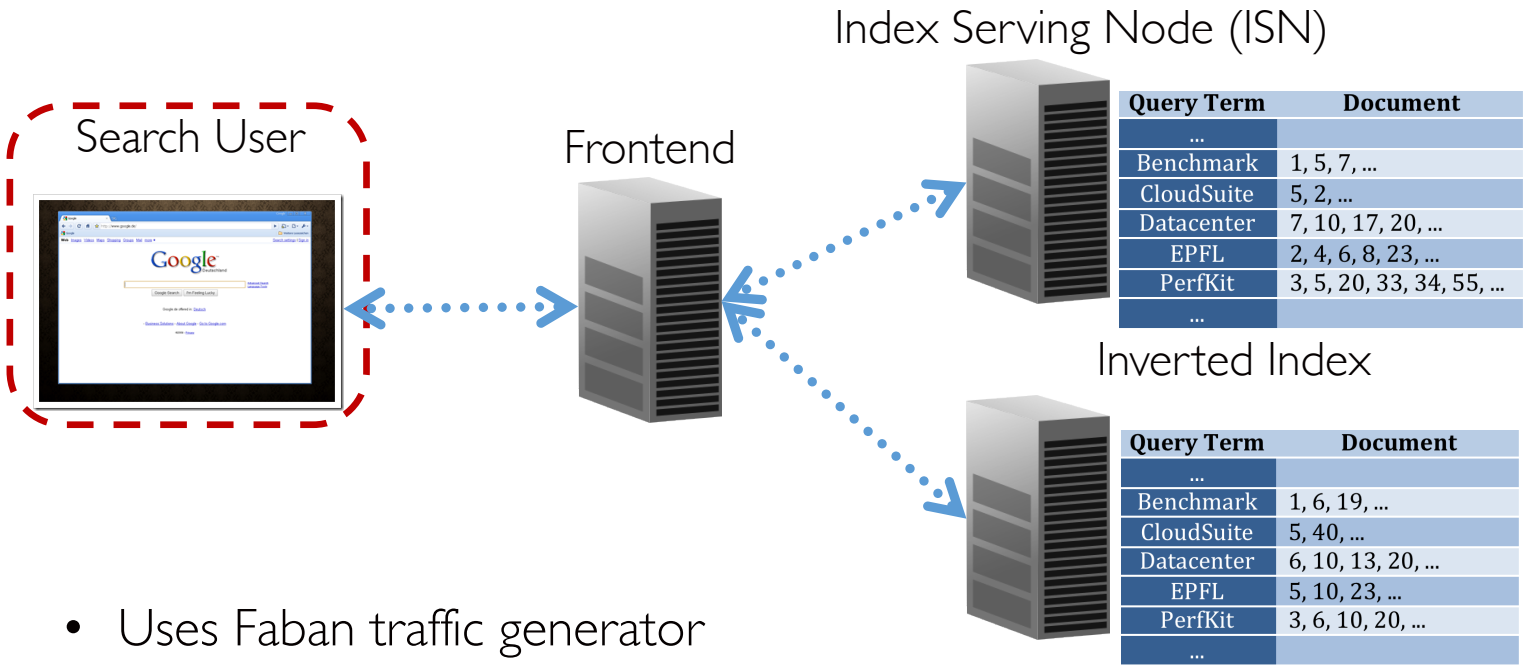

- Uses Faban traffic generator
- Flexible request mixes
  - # of terms per request from published surveys
  - Terms extracted from the crawled dataset

Web Search Operation

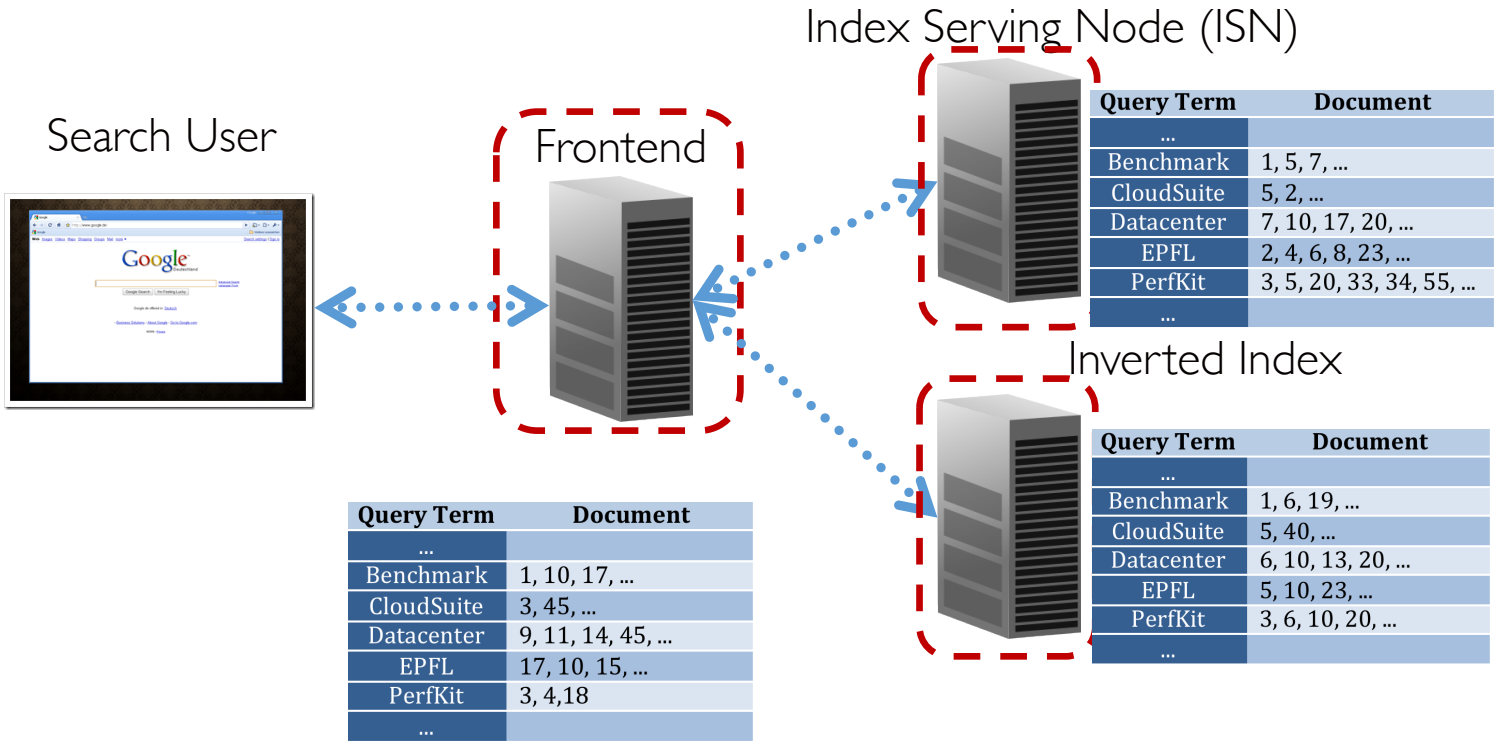

Query Term	Document
...	
Benchmark	1, 5, 7, ...
CloudSuite	5, 2, ...
Datacenter	7, 10, 17, 20, ...
EPFL	2, 4, 6, 8, 23, ...
PerfKit	3, 5, 20, 33, 34, 55, ...
...	

Query Term	Document
...	
Benchmark	1, 6, 19, ...
CloudSuite	5, 40, ...
Datacenter	6, 10, 13, 20, ...
EPFL	5, 10, 23, ...
PerfKit	3, 6, 10, 20, ...
...	

Query Term	Document
...	
Benchmark	1, 10, 17, ...
CloudSuite	3, 45, ...
Datacenter	9, 11, 14, 45, ...
EPFL	17, 10, 15, ...
PerfKit	3, 4,18
...	

- Apache Solr search engine for ISNs

Web Search Operation

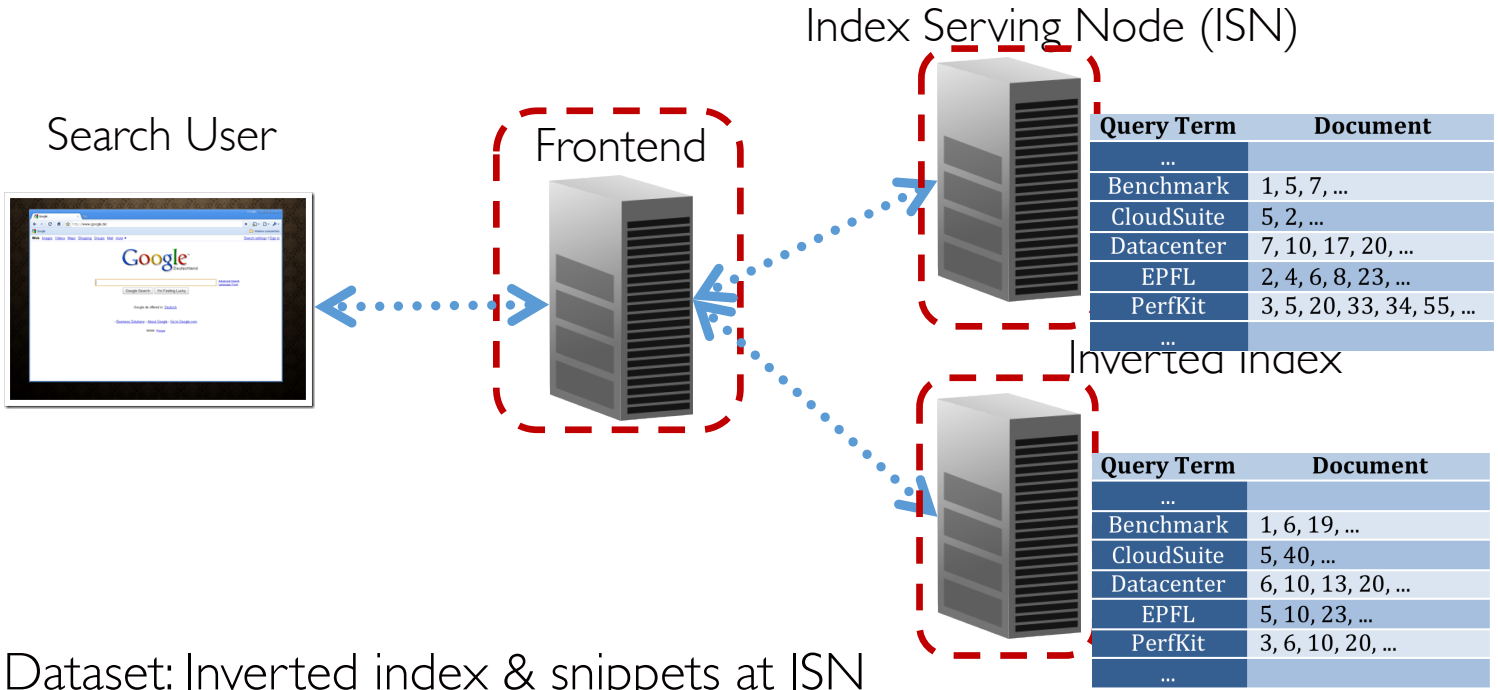

- Dataset: Inverted index & snippets at ISN
  - Generated by crawling public web (Apache Nutch)
  - Data at ISN must be memory resident
- Performance metrics: search ops/sec, N-th pct latency -

CloudSuite 3.0 Benchmarks

- Offline benchmarks
  - Data Analytics
  - Recommendation system
  - Graph Analytics
- Online benchmarks
  - Data Caching
  - Data Serving
  - Media Streaming
  - Web Search
  - Web Serving

Web Serving

- Key to all internet-based services


- All services are accessed through web servers

- Various technologies construct web content
  - HTML, PHP, JavaScript, Ruby

Web Serving Operation

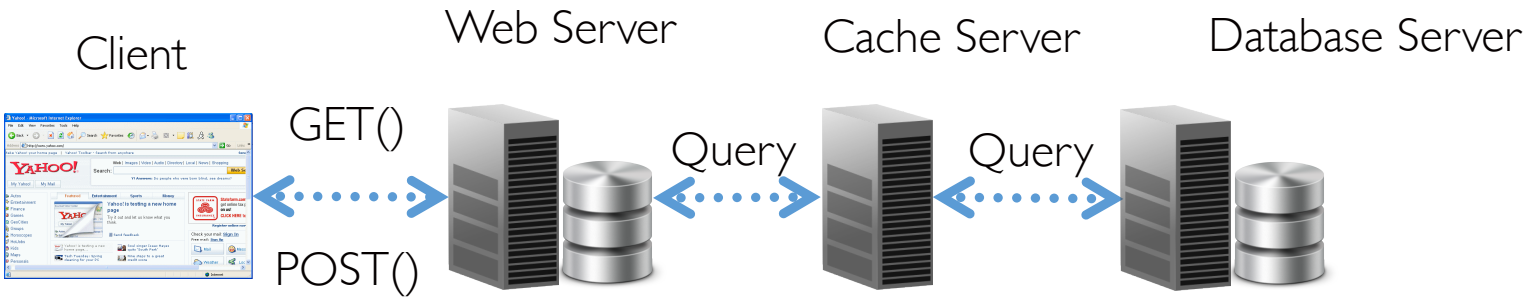

Web Serving Operation

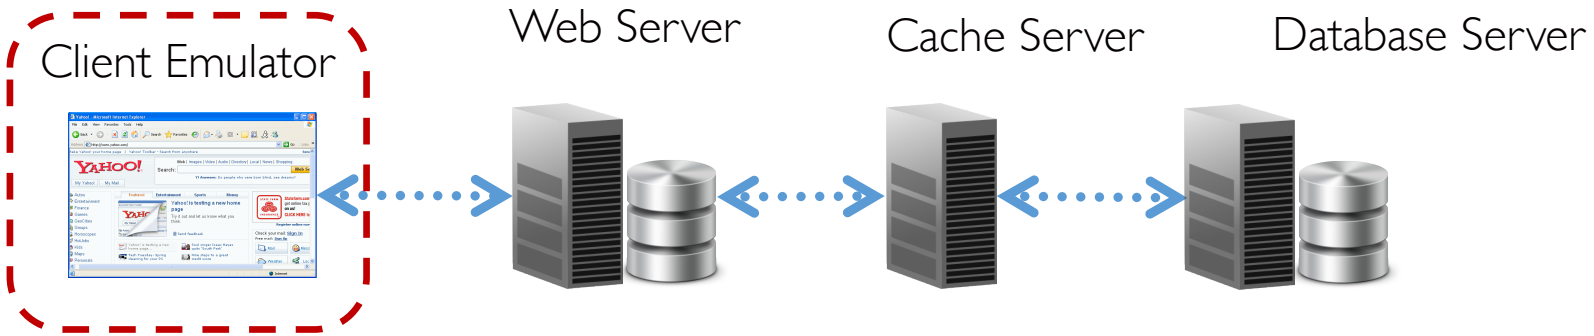

- Faban traffic generator
- Pre-configured page transition matrix (Elgg)

Web Serving Operation

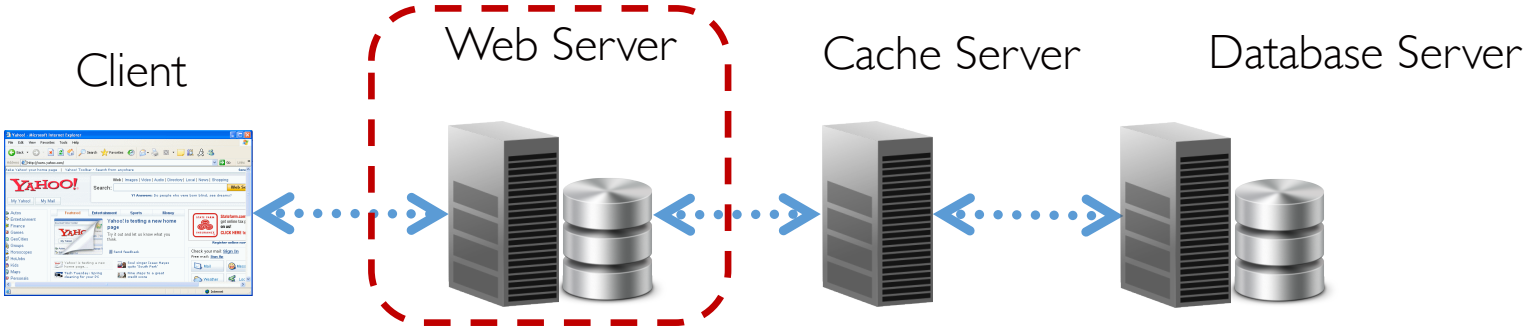

- Web server (Nginx)
- Application server (PHP)
  - Serves a social network engine (Elgg)

Web Serving Operation

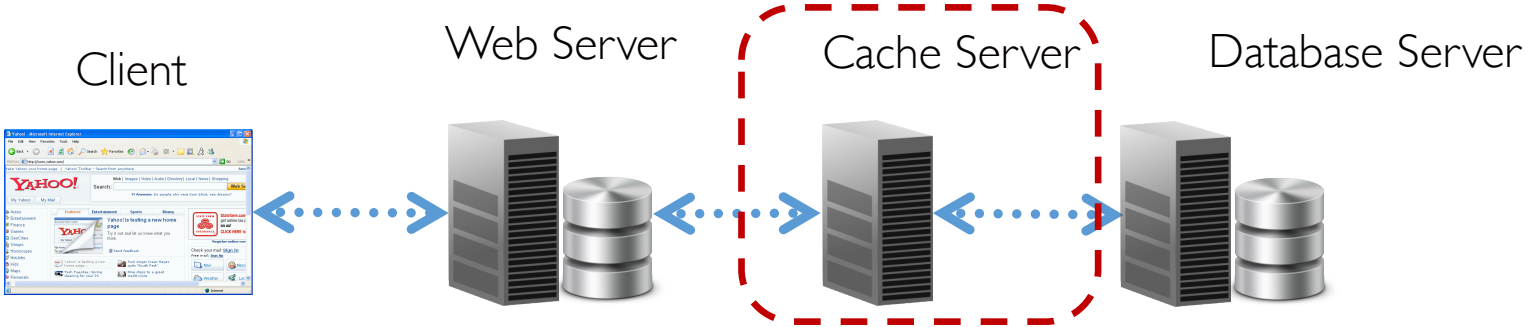

- Cache server (Memcached)

Web Serving Operation

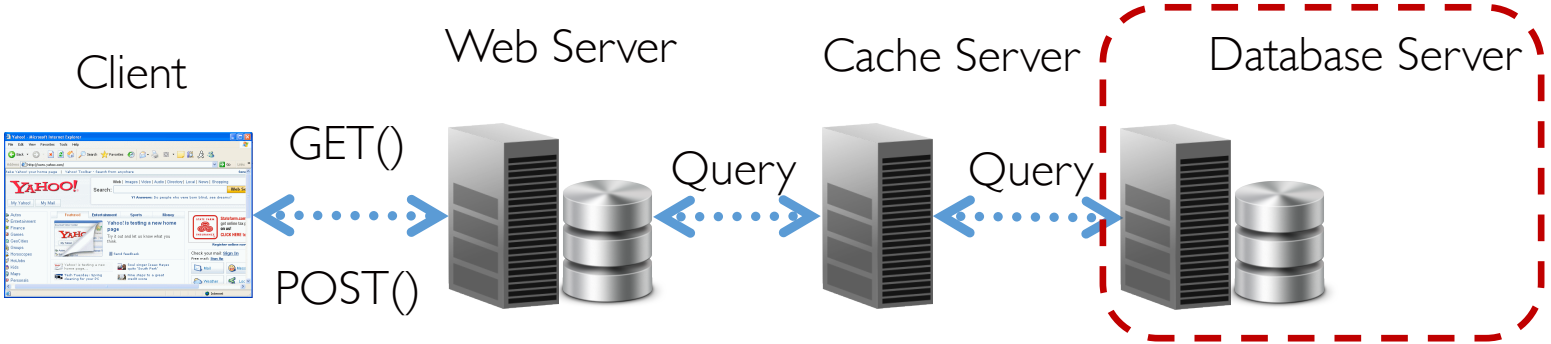

- Database server (MySQL)
- Performance metrics:
  # of pages/second served, N-th pct latency

CloudSuite 3.0 Benchmarks

- Offline benchmarks
  - Data Analytics
  - Recommendation system
  - Graph Analytics
- Online benchmarks
  - Data Caching
  - Data Serving
  - Media Streaming
  - Web Search
  - Web Serving

Future directions

- New workloads: Intelligent Personal Assistants (IPAs)
  - Examples: Apple Siri, Google Assistant, Amazon Alexa

- New open-source computer architecture simulator: QFlex
  - Stimulate architecture research on cloud workloads (e.g., CloudSuite)
  - Based on QEMU → Sharable infrastructure & reproducible results
  - Quick design space exploration and obviates for early development kits
  - Work in progress…

Conclusion

- Cloud services are becoming increasingly ubiquitous
- Current servers are built for desktop applications
- Server design needs to be driven by cloud services' needs
- CloudSuite: the benchmark suite of cloud services

Download CloudSuite 3.0

cloudsuite.ch

Thank You!

For more information please visit us at

parsa.epfl.ch