KNOWLEDGE LAB / THE UNIVERSITY OF CHICAGO

A Secure Data Enclave and Analytics Platform For Social Scientists

Yadu N. Babuji, Kyle Chard, Aaron Gerow, & Eamon Duede
Computation Institute, The University of Chicago and Argonne National Laboratory
{yadunand,chard,gerow,eduede}@uchicago.edu

2016 IEEE 12th Conference on eScience

Motivation

- Data driven research is ubiquitous. Data is fast becoming the defining assets for researchers, particularly those in the computational social sciences and humanities
- Data is increasingly large; it is also valuable, proprietary, and sensitive
- Social scientists (and other researchers) lack the technical and financial resources to securely and scalably manage large amounts of data while also supporting flexible and large-scale analytics
- Cloud computing provides “infinite” storage and compute resources, however it requires technical expertise to deploy, configure, manage, and use
- Cloud Kotta is a cloud-hosted environment that supports the secure management and analysis of large scientific datasets

With private data-sets comes great responsibility

A significant fraction of the 10TB we manage is sensitive/proprietary data

`Web of Science - from Thomson Reuters (1TB)`

`UChicago AURA grants DB - under NDA (~200GB)`

`IEEE full texts - under license (5.5TB)`

We want to make this data accessible to our colleagues and collaborators, but secured within our infrastructure.

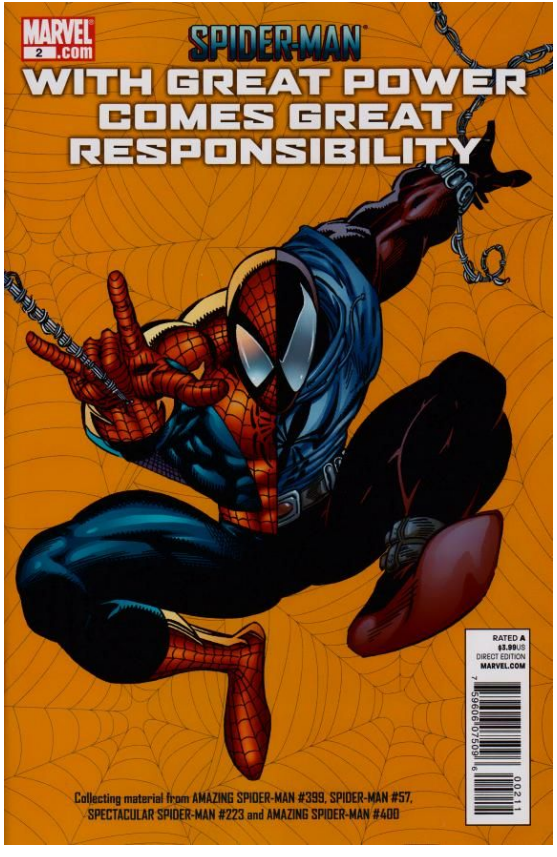

THE UNIVERSITY OF CHICAGO
KNOWLEDGE LAB

With massive data comes massive COST

We hold a tad over 10TB of research data.

10TB on EBS(SSD) = $1000 / mo

10TB on S3 (std) = $300 / mo

10TB on S3 (IA) = $125 / mo

10TB on Glacier = $70 / mo

Each comes with its own tradeoffs.

THE UNIVERSITY OF CHICAGO KNOWLEDGE LAB

Large-scale data analytics

- Analyses are user driven and often interactive
- Development is often iterative
- Analyses are often compute intensive or memory intensive
- Complex analyses can be broken down to a many-task model (SPMD) and computed in parallel
- Scientific workloads are inherently sporadic and bursty (tracking submission deadlines)
- Variable lengths of time (minutes to weeks)
- Analyses are written in many languages (e.g., Python, Julia, BaSH, C++)

With massive compute comes massive COST

We've run over **75K*** compute hours in 6 months

```
On-demand                        =    $15984.37
Spot-market (variable)           =    ~$4795.31
1 Reserved instance for 6mo      =    $17677.44
```

With i2.8xlarge, you can burn a 10K AWS credit in just 2 months.

We want to optimize for both cost and time-to-solution.

* Core hours

THE UNIVERSITY OF CHICAGO KNOWLEDGE LAB

Solution

KNOWLEDGE
LAB / THE UNIVERSITY OF CHICAGO

Cloud Kotta

- Cloud Kotta is a cloud-based platform that enables secure and cost-effective management and analysis of large, potentially sensitive data
- The platform automatically provisions cloud infrastructure to host user submitted jobs
- Data is migrated between storage tiers depending on access patterns and pre-defined policies
- Role based access model for security

in Malayalam Kotta means Fortress

THE UNIVERSITY OF CHICAGO KNOWLEDGE LAB

* Pictured: Mehrangarh Fort at Jodhpur, Rajasthan

Automated storage management

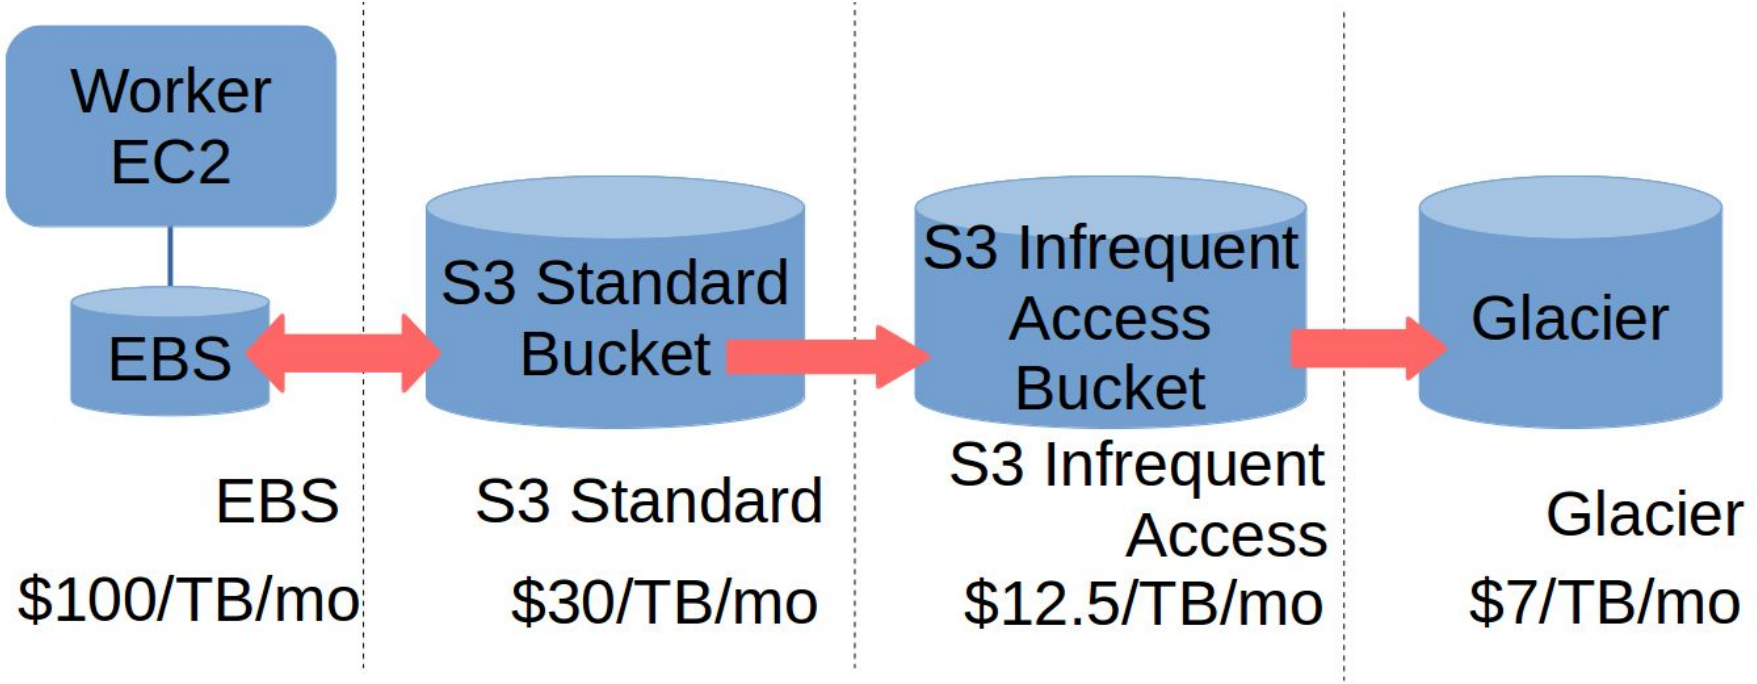

Elastic Provisioning

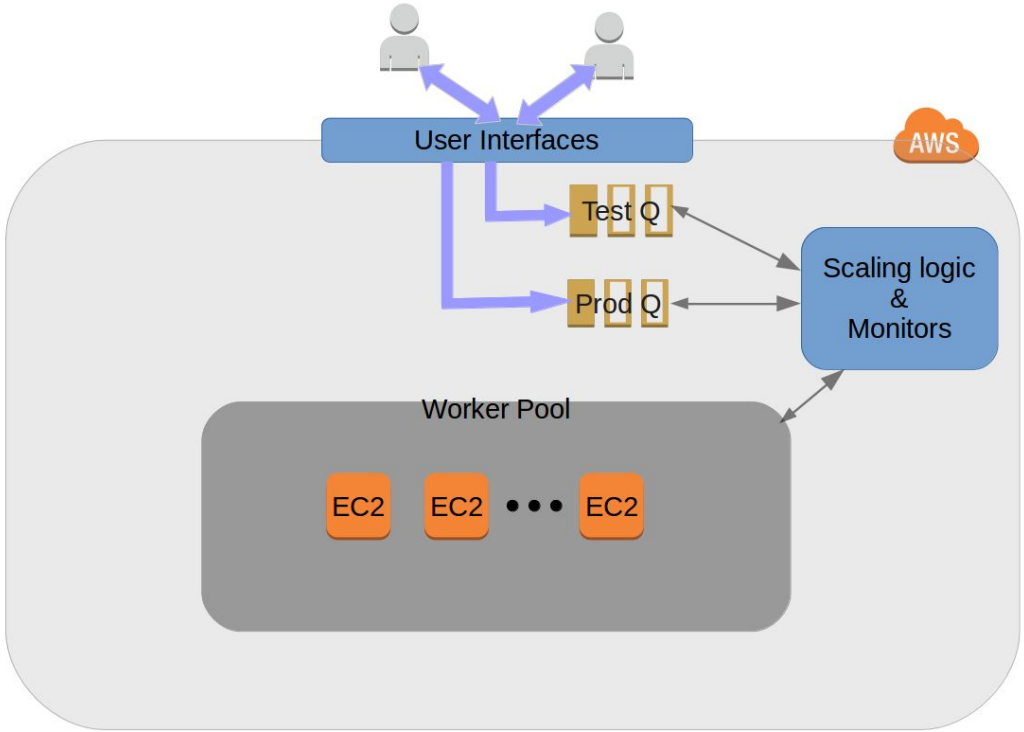

Security model

- Principle of least privilege throughout
- “Log in with Amazon”
- Users are assigned roles
- Policies permit access to resources for individual roles
- Instances are granted a trusted role that allows them to switch to a user role temporarily in order to inherit user permissions (e.g., access secure data)
- Compute layer is hosted within a private subnet enclosed within a VPC

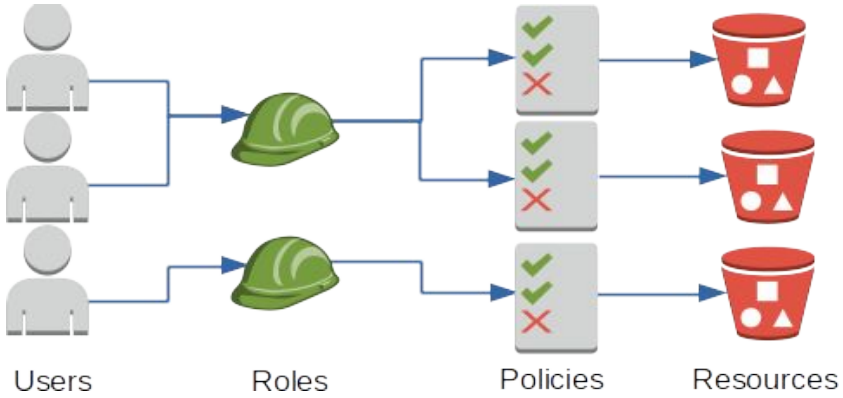

Cloud Formation

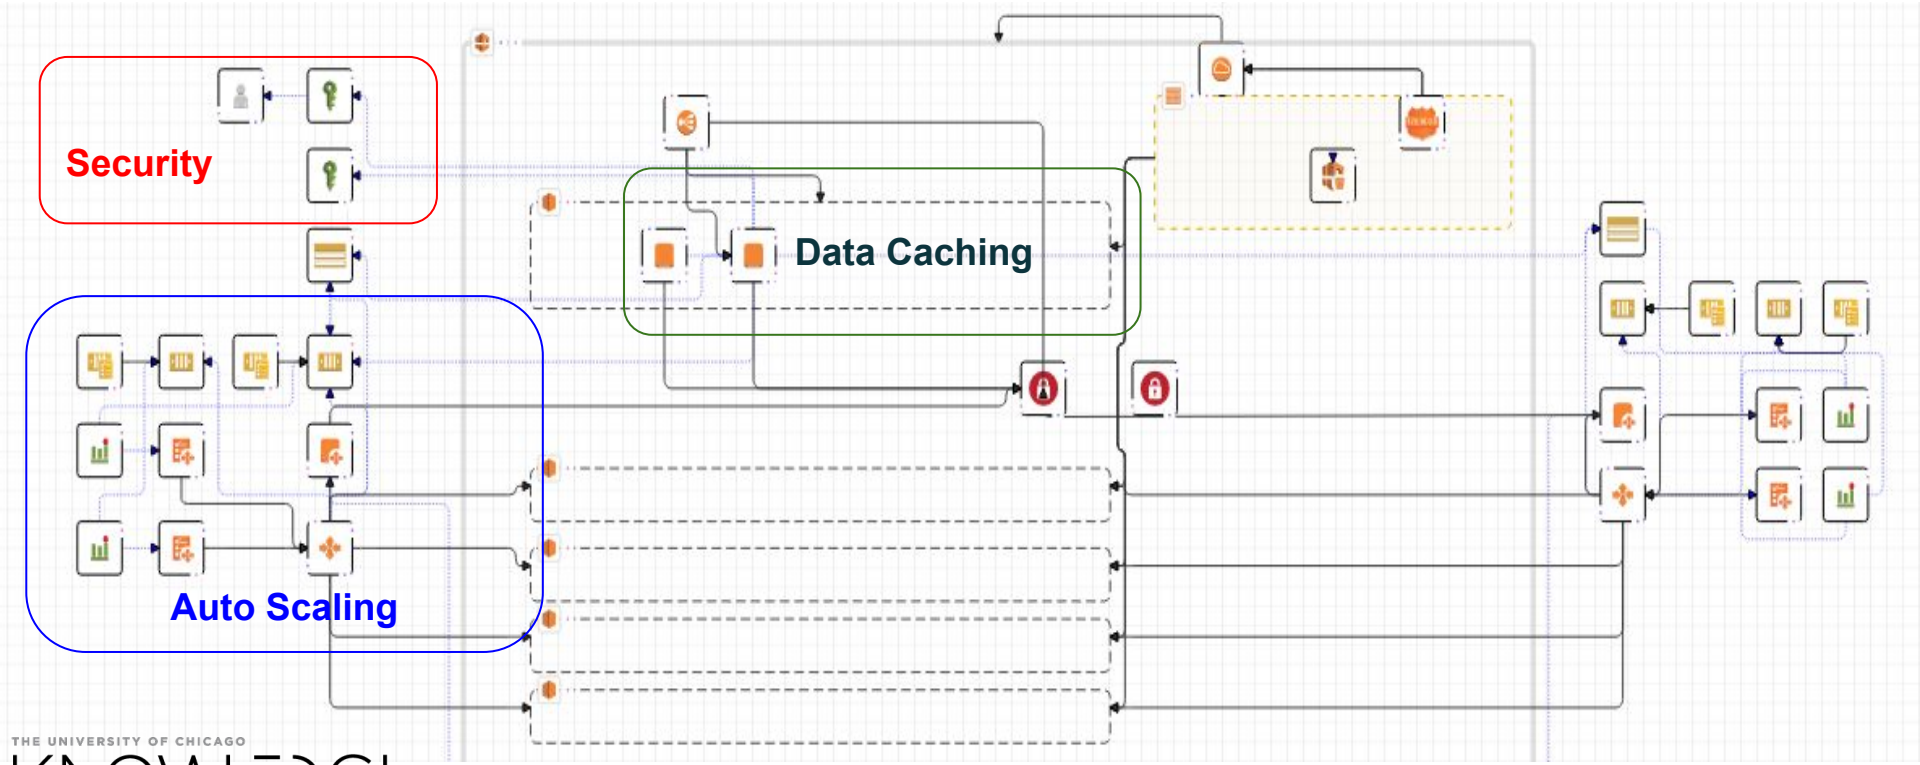

User Interfaces

Web Interface

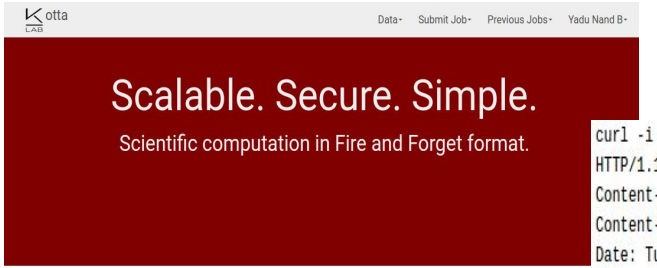

REST API

```
curl -i -H "Accept: application/json" http://52.2.217.165:8888/rest/v1/status
HTTP/1.1 200 OK
Content-Length: 3392
Content-Type: application/json
Date: Tue, 25 Oct 2016 05:43:38 GMT
Server: ip-10-0-0-93

{"status": "completed",
 "items":
   {"0": {"username": "Yadu Nand B"},
    "1": {"z_stagein_dur": 0.041018962860107421875},
    "2": {"executable": "/bin/bash myscript.sh"},
    "3": {"job_id": "d782d28d-4407-44df-8aec-a128a64427f9"},
    "4": {"complete_time": "2016-10-20 18:29:07"},
    "5": {"submit_stamp": "2016-10-20 18:32:56"},
    "6": {"z_processing_dur": -0.9968240261077880859375},
    "7": {"status": "completed"}, "8": {"outputs": "<a href=\"https://klab-jo
```

Command Line Interface

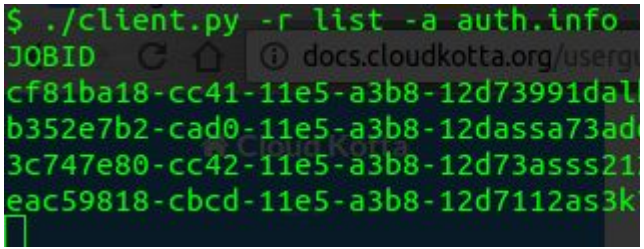

THE UNIVERSITY OF CHICAGO
KNOWLEDGE LAB

User Workflow

Data Interface

Upload Data

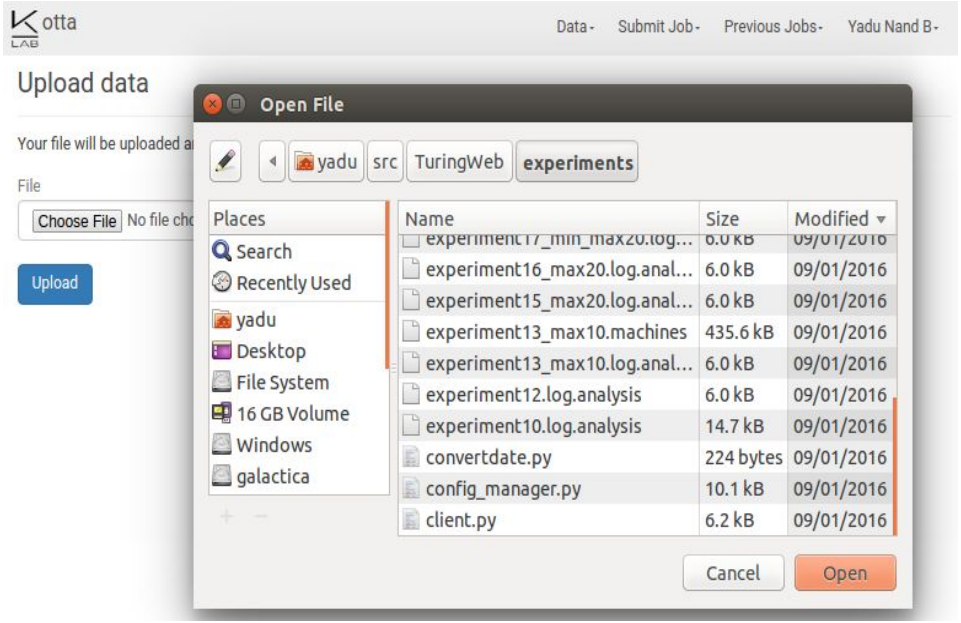

Browse Data

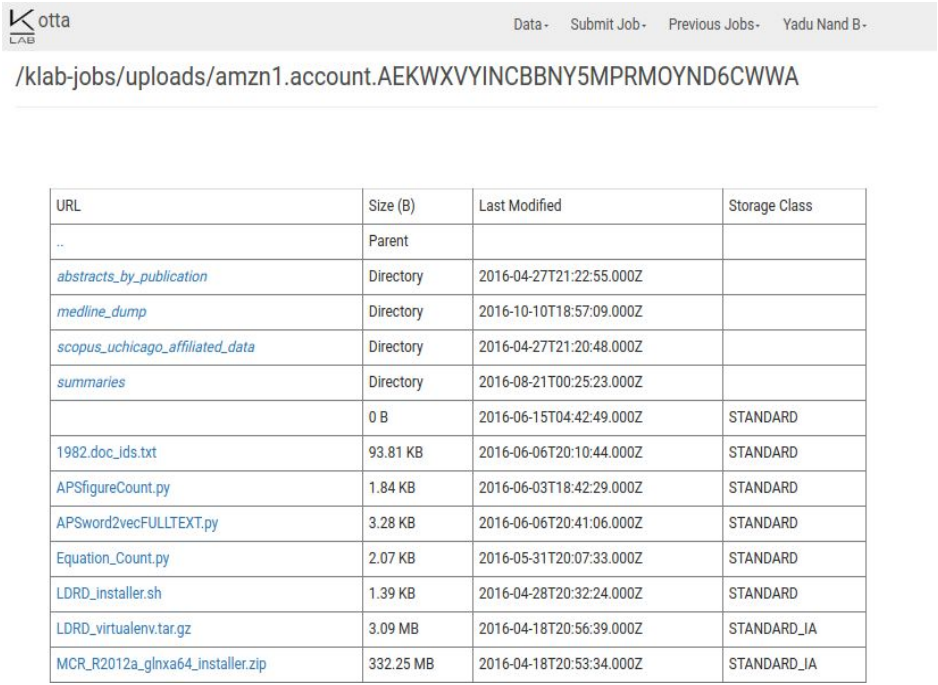

URL	Size (B)	Last Modified	Storage Class
..	Parent		
abstracts_by_publication	Directory	2016-04-27T21:22:55.000Z	
medline_dump	Directory	2016-10-10T18:57:09.000Z	
scopus_uchicago_affiliated_data	Directory	2016-04-27T21:20:48.000Z	
summaries	Directory	2016-08-21T00:25:23.000Z	
	0 B	2016-06-15T04:42:49.000Z	STANDARD
1982.doc_ids.txt	93.81 KB	2016-06-06T20:10:44.000Z	STANDARD
APSfigureCount.py	1.84 KB	2016-06-03T18:42:29.000Z	STANDARD
APSword2vecFULLTEXT.py	3.28 KB	2016-06-06T20:41:06.000Z	STANDARD
Equation_Count.py	2.07 KB	2016-05-31T20:07:33.000Z	STANDARD
LDRD_installer.sh	1.39 KB	2016-04-28T20:32:24.000Z	STANDARD
LDRD_virtualenv.tar.gz	3.09 MB	2016-04-18T20:56:39.000Z	STANDARD_IA
MCR_R2012a_glnxa64_installer.zip	332.25 MB	2016-04-18T20:53:34.000Z	STANDARD_IA

THE UNIVERSITY OF CHICAGO
KNOWLEDGE LAB

Job Submission

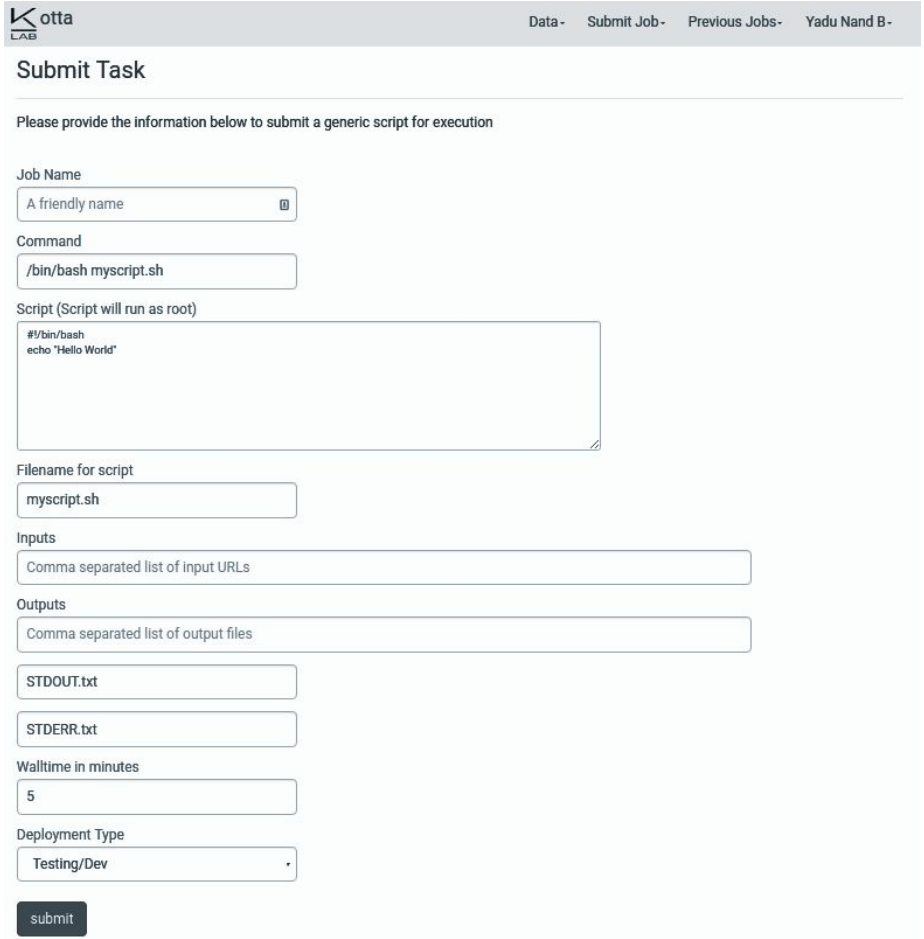

THE UNIVERSITY OF CHICAGO KNOWLEDGE LAB

Job management

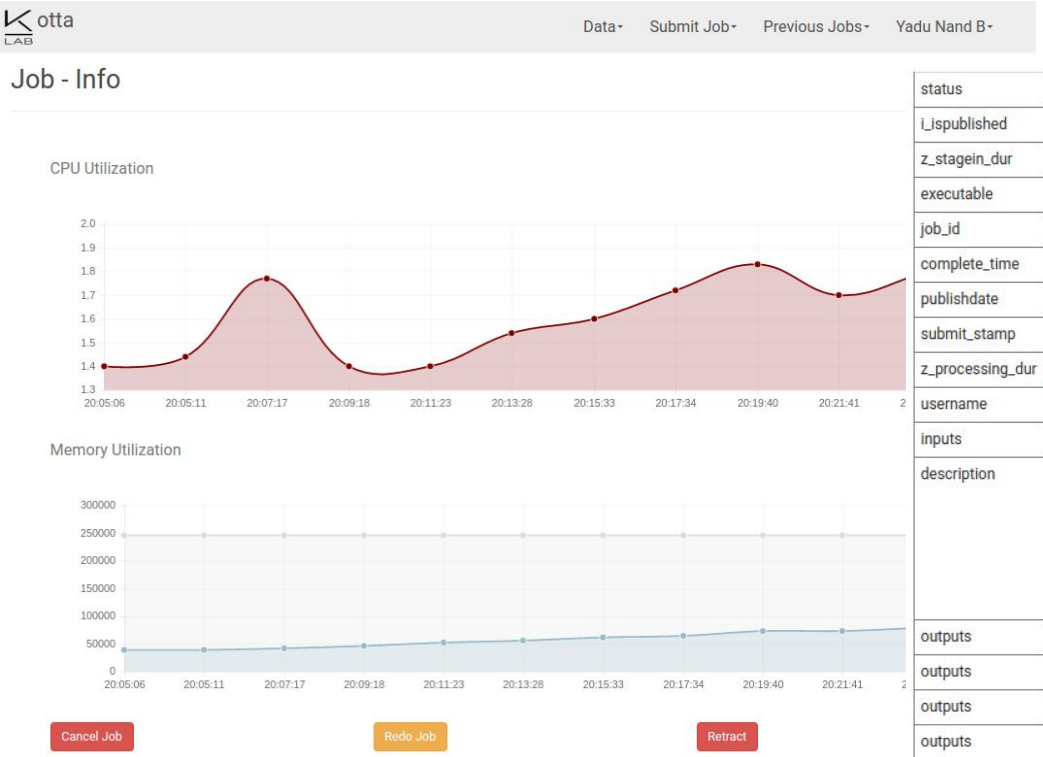

status	completed
i_ispublished	1
z_stagein_dur	329.211119174957275390625
executable	/bin/bash myscript.sh
job_id	0872b9d1-38cb-413c-a5d3-8a9a9a2de818
complete_time	2016-07-11 20:26:44
publishdate	2016-07-13 16:19:22
submit_stamp	2016-07-11 19:54:00
z_processing_dur	1295.0844790935516357421875
username	Wanqi Zhu
inputs	article_influence.py
description	Job Behavior: Computes the Article Influence of journals based on eigenfactor and article counts for # of papers published. Data obtained from web of science. Done in parallel insanely fast, but it's not necessary as it's not computationally intensive. Inputs: journal_counts_1995.tsv ~ journal_counts_2015.tsv, one per year, in tab-separated values of (including headers for the first row) JOURNAL COUNT eigenfactors_1995.tsv ~ eigenfactors_2015.tsv, one per year, in tab-separated values of (no headers) JOURNAL EIGENFACTOR Outputs: Calculate the Article Influence for each of the years and output to article_influence.csv. Also outputs intermediate values for debugging, but they are not useful: t_counts.csv, j_counts.csv, j_ef.csv
outputs	myscript.sh
outputs	article_influence.csv
outputs	t_counts.csv
outputs	j_counts.csv
outputs	j_ef.csv
outputs	STDERR.txt
outputs	STDOUT.txt
start_time	2016-07-11 19:59:37

THE UNIVERSITY OF CHICAGO
KNOWLEDGE LAB

Early Usage/Results

System Utilization

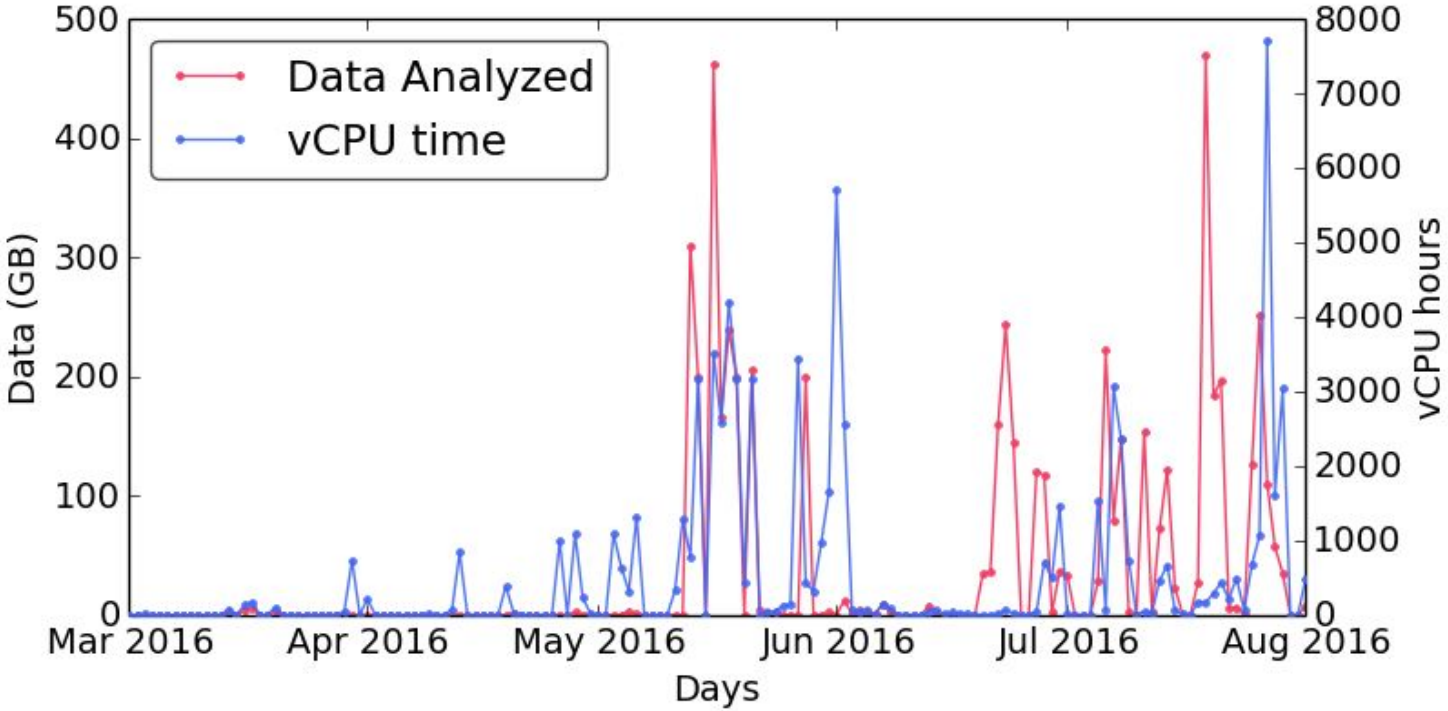

Elastic scaling experiment

- To demonstrate the automatic scaling behavior we used a test-workload derived from historical production usage
- 40 jobs of 1,3, or 4 hour durations with inter-arrival time from poisson-distribution(λ = 0.1667).
- Jobs simply call `sleep()`
- Each job uses a randomly selected data input of size {1,3,5,7,9}GB
- The scaling limit was set to a maximum of 40 nodes
- We plot the total nodes active and idle, as well as the state of each of the 40 jobs. X axis is time.

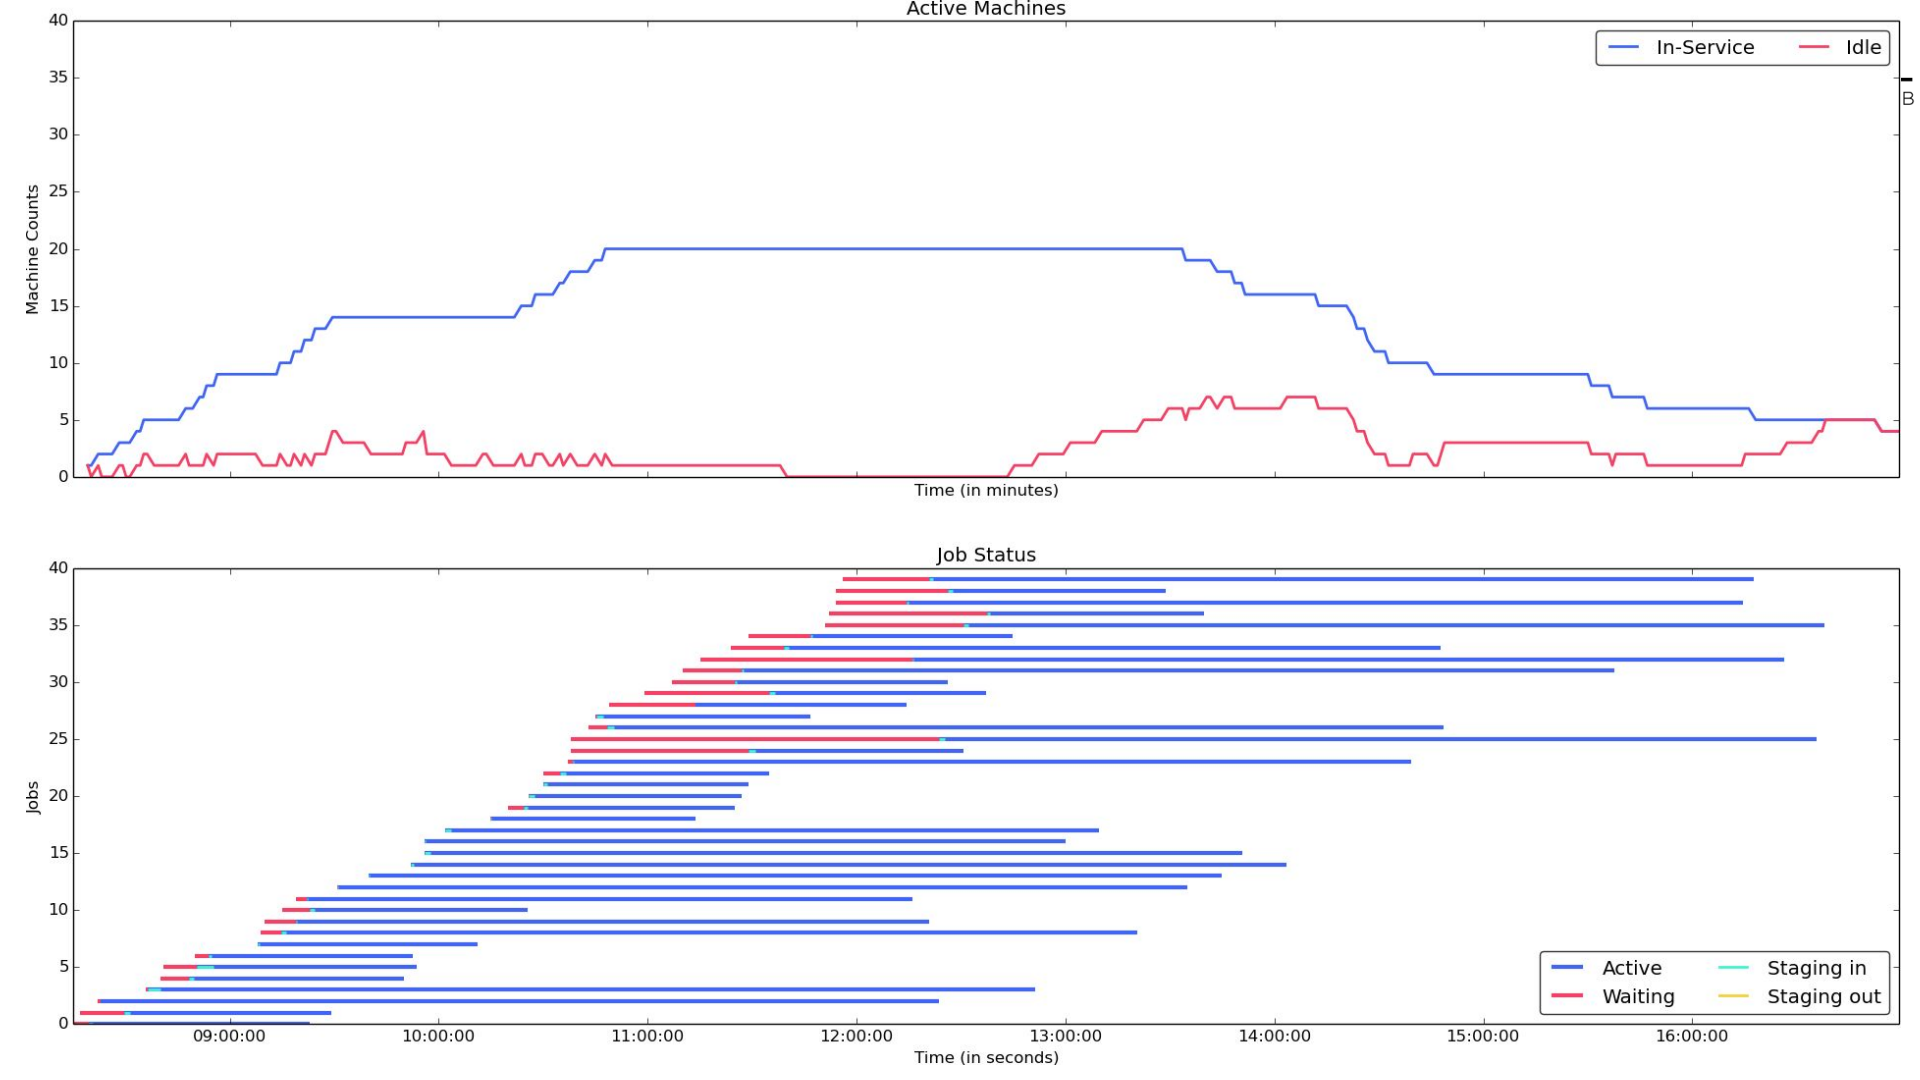
Active Machines
In-Service
Idle
Machine Counts
Time (in minutes)
Job Status
Active
Waiting
Staging in
Staging out
Jobs
09:00:00
10:00:00
11:00:00
12:00:00
13:00:00
14:00:00
15:00:00
16:00:00
Time (in seconds)

Early science on Cloud Kotta

- Text Analytics
- Matrix Factorization
- Optical Character Recognition (tesseract)
- Network Analysis
- Author-Topic models

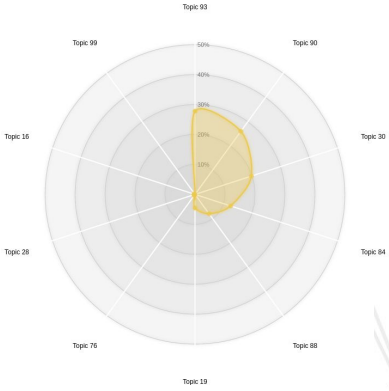

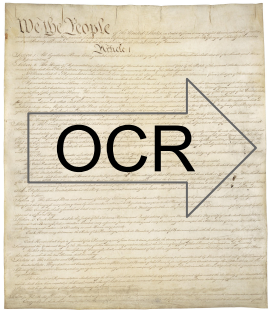

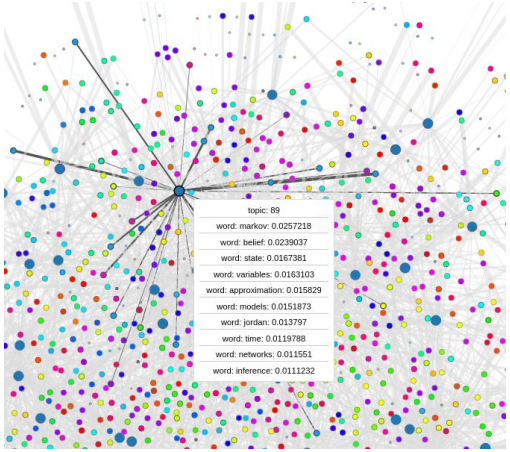

THE UNIVERSITY OF CHICAGO
KNOWLEDGE LAB

Acknowledgements

KNOWLEDGE LAB / THE UNIVERSITY OF CHICAGO

IBM

John Templeton Foundation

facebook

Thanks

- Github repo : **https://github.com/yadudoc/cloud_kotta**
- Documentation : **http://docs.cloudkotta.org/**
- Support : **yadunand@uchicago.edu**

THE UNIVERSITY OF CHICAGO
KNOWLEDGE LAB