

# Cloud review, activities, evolution

Peter Love  
21 September 2017



# Introduction

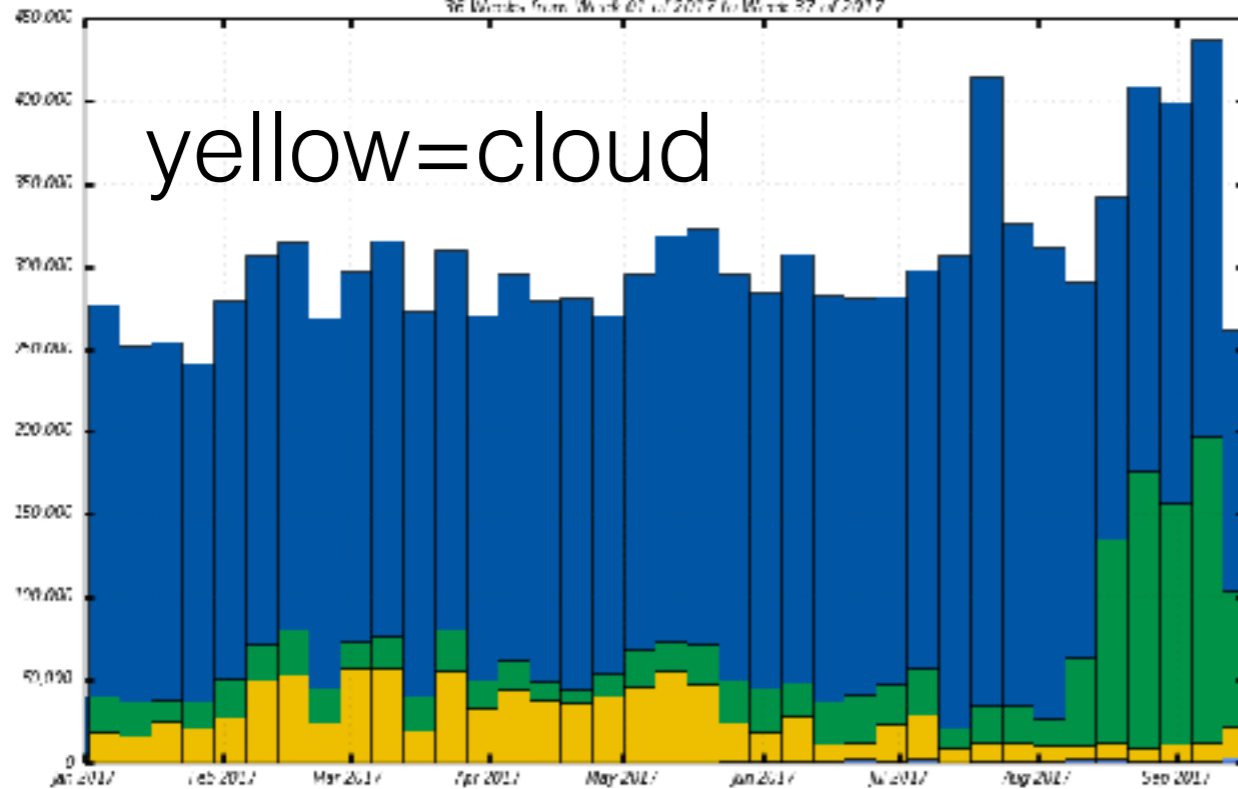
- Cloud compute usage review
- Objectstore developments
- Growth of commercial services
- Forward looking and opportunities for ATLAS cloud activity

# Usage 2017 calendar



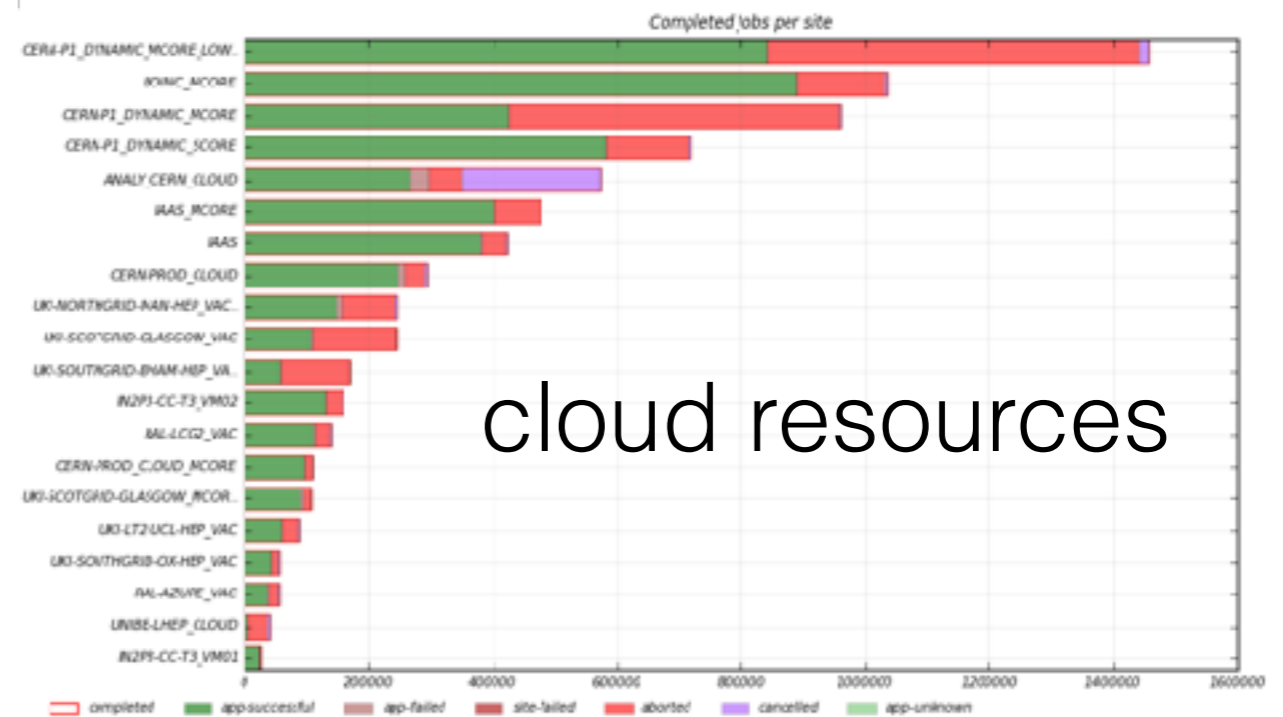
Slots of Running Jobs

36 Weeks from Week 01 of 2017 to Week 37 of 2017




■ grid ■ fpc ■ cloud ■ local ■ None

Maximum: 1,171,112, Minimum: 0,000, Average: 269,194, Current: 261,257



cloud resources

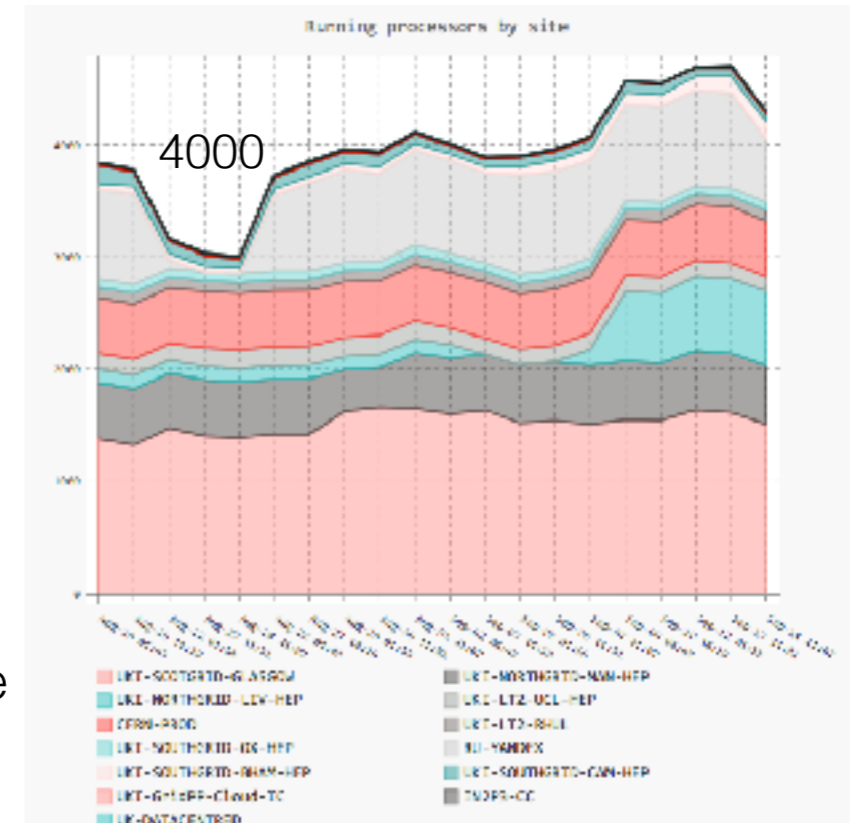
Completed Jobs per site



■ completed ■ app successful ■ app failed ■ site failed ■ aborted ■ cancelled ■ app unknown

# Active resources

- Dominated by CERN (~1000) and IAAS (~5000) (7000 split between ATLAS/Belle)
- Growing number of Vac sites with GridPP consolidation of sites (ATLAS Vac ~1000 slots).
- Cloudscheduler instances at CERN/uVic
- IN2P3-CC-T3 testing Openstack 'Synergy', which allows fairshare and consequently opportunistic use of idle resources



**IAAS** 13:17:26 15-Sep

amazon-east (20) amazon-west (20) beaver (21) cc-east (10) cc-west-hep (330) charonion (20) dal-hb (6) dal-go (6) tr (2)

Cloud	Starting	CloudScheduler VMs					HTCondor Slots										HTCondor Jobs					
		Running	Retiring	Error	Idle	Lost	Unregistered	1	2	3	4	5	6	7	8	Held	Total	Idle	Running	Completed	Held	
beaver	0	21	0	0	0	0	21	0	0	0	0	0	0	0	0	0	0	468	98	365	5	0
cc-west-hep	6	267	77	0	0	6	344	0	0	0	0	0	0	0	0	0	0	465	98	365	2	0

Jobs	Total	Idle	Running	Completed	Held
All	468	98	365	5	0
1 Core	3	0	0	3	0
8 Core	465	98	365	2	0

CS System Status						
Load	Cpu	Ram	Swap	SI	So	Disk
~	~	~	~	~	~	41

**ATLAS-Cern** 13:15:48 15-Sep

cern-atlas (100) cern-victoria (4) datacentro (2) gridpp-imprial (2) gridpp-oxford (10) tr (10) rooter (10)

Cloud	Starting	CloudScheduler VMs					HTCondor Slots										HTCondor Jobs					
		Running	Retiring	Error	Idle	Lost	Unregistered	1	2	3	4	5	6	7	8	Held	Total	Idle	Running	Completed	Held	
cern-atlas	0	100	0	0	0	0	100	0	0	0	0	0	0	0	0	0	148	31	117	0	0	
cern-victoria	0	1	0	0	0	0	1	0	0	0	0	0	0	0	0	0	130	18	112	0	0	
cern-extension	0	4	0	0	0	0	4	0	0	0	0	0	0	0	0	0						
tr	0	12	0	0	0	0	12	0	0	0	0	0	0	0	0	0	4	3	88	1	0	31

CS System Status						
Load	Cpu	Ram	Swap	SI	So	Disk
~	~	~	~	~	~	31

# Storage developments

- Recent pre-GDB covered objectstore utility
- Integration of objectstores
  - Enabling usage via Dynafed advanced and meeting problems with existing models, eg. handling of replicas (see uVic presentation)
  - RAL's site in Azure tested using Dynafed to handle authz/authn exploiting S3/Swift APIs on ECHO (Ceph)
- Step back from using objectstore with ATLAS event service
- Performance measurements and integration with Hammercloud monitoring are underway

# Cloud services beyond compute and storage

- Growth of cloud compute resources is slow, at least resources available to ATLAS
- Adoption of container technology within sites, not yet container cluster hosting. eg. from recent GridPP meeting
  - RAL compute now entirely built using HTCondor Docker universe, 22000 cores with services not just WNs
  - Glasgow migrating compute to container fabric using Vac model
- Besides compute, what commercial services are available? Take a survey of cloud based web services



Compute



Storage



Database



Migration



Networking & Content Delivery

Amazon EC2

Virtual Servers in the Cloud

Amazon EC2 Container Registry

Store and Retrieve Docker Images

Amazon EC2 Container Service

Run and Manage Docker Containers

Amazon Lightsail

Launch and Manage Virtual Private Servers

Amazon VPC

Isolated Cloud Resources

AWS Batch

Run Batch Jobs at Any Scale

AWS Elastic Beanstalk

Run and Manage Web Apps

AWS Lambda

Run Your Code in Response to Events

Auto Scaling

Automatic Elasticity



Developer Tools



Management Tools



Security, Identity & Compliance



Analytics



Artificial Intelligence



Mobile Services



Application Services



Messaging



Business Productivity



Desktop & App Streaming




Internet of Things



Contact Center



Game Development



Compute



Storage



Database



Migration



Networking & Content Delivery

Amazon S3

Scalable Storage in the Cloud

Amazon EBS

Block Storage for EC2

Amazon Elastic File System

Managed File Storage for EC2

Amazon Glacier

Low-Cost Archive Storage in the Cloud

AWS Storage Gateway

Hybrid Storage Integration

AWS Snowball

Petabyte-scale Data Transport

AWS Snowball Edge

Petabyte-scale Data Transport with On-board Compute

AWS Snowmobile

Exabyte-scale Data Transport



Developer Tools



Management Tools



Security, Identity & Compliance



Analytics



Artificial Intelligence



Mobile Services



Application Services



Messaging



Business Productivity



Desktop & App Streaming



Internet of Things



Contact Center



Game Development





Compute



Storage



Database



Migration



Networking & Content Delivery

### Amazon Aurora

High Performance Managed Relational Database

### Amazon RDS

Managed Relational Database Service for MySQL, PostgreSQL, Oracle, SQL Server, and MariaDB

### Amazon DynamoDB

Managed NoSQL Database

### Amazon DynamoDB Accelerator (DAX)

Fully Managed In-Memory Cache for DynamoDB

### Amazon ElastiCache

In-Memory Caching System

### Amazon Redshift

Fast, Simple, Cost-Effective Data Warehousing

### AWS Database Migration Service

Migrate Databases with Minimal Downtime



Developer Tools



Management Tools



Security, Identity & Compliance



Analytics



Artificial Intelligence



Mobile Services



Application Services



Messaging



Business Productivity



Desktop & App Streaming




Internet of Things



Contact Center



Game Development



Compute



Storage



Database



Migration



Networking & Content Delivery



Developer Tools



Management Tools



Security, Identity & Compliance



Analytics



Artificial Intelligence

**Amazon Athena**

Query Data in S3 using SQL

**Amazon EMR**

Hosted Hadoop Framework

**Amazon CloudSearch**

Managed Search Service

**Amazon Elasticsearch Service**

Run and Scale Elasticsearch Clusters

**Amazon Kinesis**

Work with Real-Time Streaming Data

**Amazon Redshift**

Fast, Simple, Cost-Effective Data Warehousing

**Amazon QuickSight**

Fast Business Analytics Service

**AWS Data Pipeline**

Orchestration Service for Periodic, Data-Driven Workflows

**AWS Glue**

Prepare and Load Data



Mobile Services



Application Services



Messaging



Business Productivity



Desktop & App Streaming



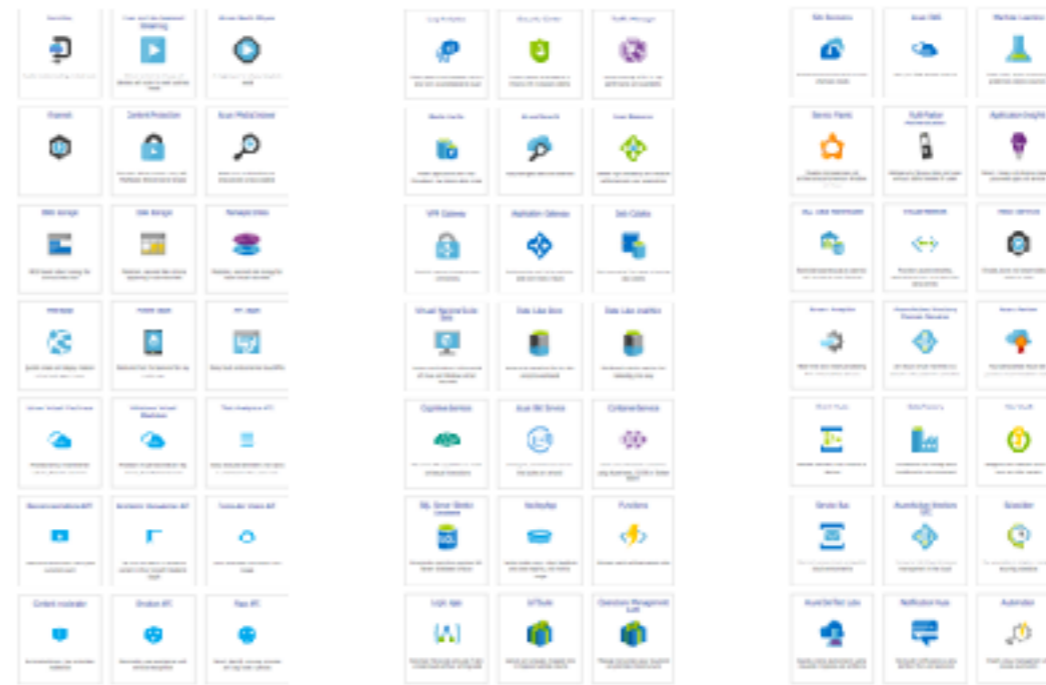
Internet of Things



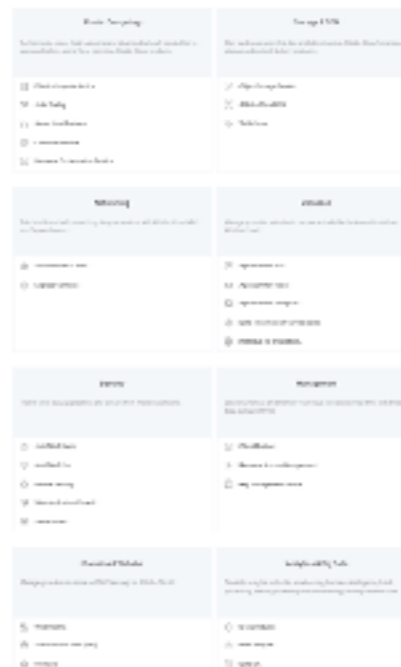
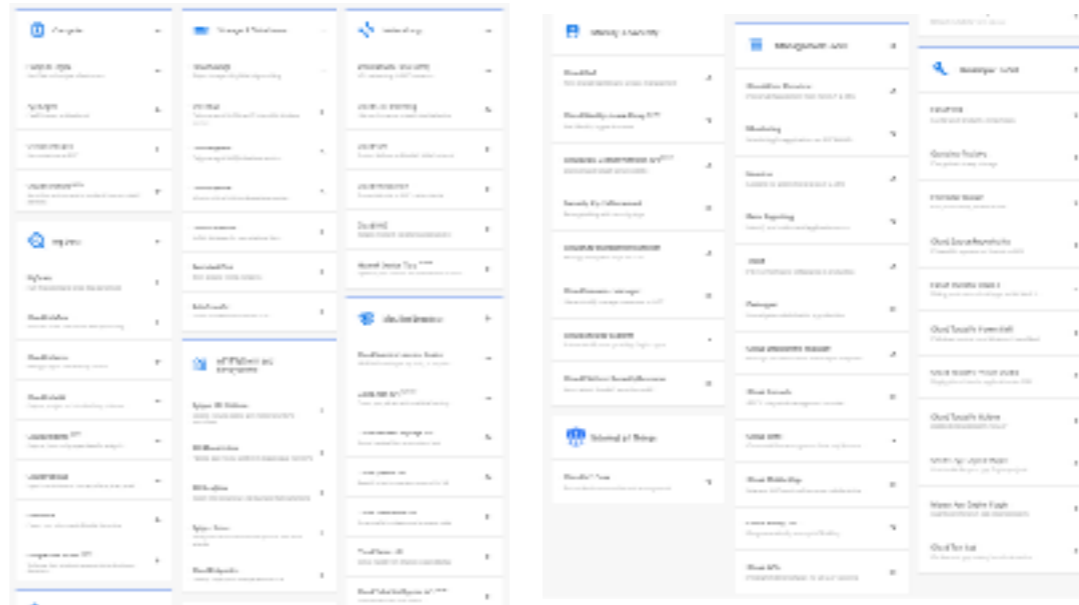
Contact Center



Game Development



Google Cloud Platform



Similar story to AWS

# Commercial services

- Very diverse set of services providing a rich ecosystem, depending on the provider
- Try these systems, evaluate cost & utility
- Feedback to ATLAS/CERN private services and help guide ATLAS requirements  
(See 'ADC themes' - Torre)

# Public and private services

- Example of public services used by ATLAS (besides compute/storage)
  - AWS DynamoDB (serverless DB)
  - AWS Athena (confusingly named for us, this is SQL on S3 objects)
  - Developer tools to smooth workflows, again, feedback into ATLAS private solutions
- Example of services deployed privately, reasonably mature
  - CERN DB on demand
  - CERN container hosting - Openstack Magnum (k8,swarm,mesos)
  - CERN streaming analytics platform - Hadoop, Spark, Kafka
  - CERN messaging/logging services - ELK
  - MWT2 elasticsearch service

Currently hosting  
-- 380 MySQL  
-- 110 PostgreSQL  
-- 60 InfluxDB  
-- 10 Oracle (Phasing out)

- Our goals: **service** delivery and working on selected **projects** with the user community

- GDB, Luca Canali

# An example of AWS-Athena (unfortunate name for us)

## Amazon Athena

Interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL



Query S3 directly  
and right away



No need to move  
data to Athena



No infrastructure  
or clusters to set  
up and manage



Results within  
seconds




Pay for just the  
queries you run

Dramatically lowers the barrier to entry

# AWS-Athena example usecase

- Pilot factory job records stored in S3, daily uploads using 'bot' write-only credentials, one-time setup
- AWS-Athena table definition
- SQL queries directly on S3 data, serverless
- Query cost based on amount scanned
- This case \$0.05 per full 1TB scan
- 1TB sample, S3 dominates \$2/day cost





▼ All services

- Compute
  - EC2
  - EC2 Container Service
  - Lightsail
  - Elastic Beanstalk
  - Lambda
  - Batch

- Storage
  - S3
  - EFS
  - Glacier
  - Storage Gateway

- Database
  - RDS
  - DynamoDB
  - ElastiCache
  - Amazon Redshift

- Networking & Content Delivery
  - VPC
  - CloudFront
  - Direct Connect
  - Route 53

- Migration
  - AWS Migration Hub
  - Application Discovery Service
  - Database Migration Service
  - Server Migration Service
  - Snowball

- Developer Tools
  - CodeStar
  - CodeCommit
  - CodeBuild
  - CodeDeploy
  - CodePipeline
  - X-Ray

- Management Tools
  - CloudWatch
  - CloudFormation
  - CloudTrail
  - Config
  - OpsWorks
  - Service Catalog
  - Trusted Advisor
  - Managed Services

- Security, Identity & Compliance
  - IAM
  - Inspector
  - Certificate Manager
  - Directory Service
  - WAF & Shield
  - Artifact
  - Amazon Macie
  - CloudHSM

- Analytics
  - Athena
  - EMR
  - CloudSearch
  - Elasticsearch Service
  - Kinesis
  - Data Pipeline
  - QuickSight
  - AWS Glue

- Artificial Intelligence
  - Lex
  - Amazon Polly
  - Rekognition
  - Machine Learning

- Internet of Things
  - AWS IoT
  - AWS Greengrass

- Contact Center
  - Amazon Connect

- Game Development
  - Amazon GameLift

- Mobile Services
  - Mobile Hub
  - Cognito
  - Device Farm
  - Mobile Analytics
  - Pinpoint

- Application Services
  - Step Functions
  - SWF
  - API Gateway
  - Elastic Transcoder

- Messaging
  - Simple Queue Service
  - Simple Notification Service
  - Simple Email Service

- Business Productivity
  - WorkDocs
  - WorkMail
  - Amazon Chime

- Desktop & App Streaming
  - WorkSpaces
  - AppStream 2.0

Used

Useful?



# Summary

- Cloud compute usage is stable but not growing
- Compute via container cluster hosting on the horizon
- Objectstore developments continue and a few interesting usecases are being studied
- Commercial services have diverse set of mature services, providing a *integrated stack* of services
- Similar services deployed privately currently used by ATLAS, explore and grow these

# Discussion points

- Should we bother to expose these services?
- If so how to identify usecases?
- Example: how are machine learning platforms for users? Easy to start using?
- Can we see something like DAST assistance from CloudSearch, Amazon LEX chatbot anyone?
- What about more pragmatic options like Lambda, Batch, Kubeless?
- Integrated 'Management tools' are far more advanced than the equivalent private interfaces, particularly Identity, Accounting, and Networking