AAM NATIONAL CAMPAIGN TECH TALK: INTRODUCTION TO AMAZON WEB SERVICES (AWS)

Irene Smith, Software Engineer

January 21st, 2021

Intro to Tech Talks

- Purpose of these talks is to engage with the community on types of technologies we are using and developing, there are many more planned
- Ground rules for talk
  - Answers to questions you have may be in upcoming slides
  - It's okay to ask an important question on a slide, but if it can wait then please do so
  - Mute your mike unless you need to talk
  - We'll keep an issues Parking Lot to keep the Tech Talk on point and on time
  - Remember the Tech Talk is being recorded for NASA and its Partners
- Recording. We are recording these Tech Talks and will post online, once approved for external release

Intro to Amazon Web Services (AWS)

- What is Cloud Computing?
- AWS Basics
- Cloud Computing within NASA
- AOM's Use of AWS

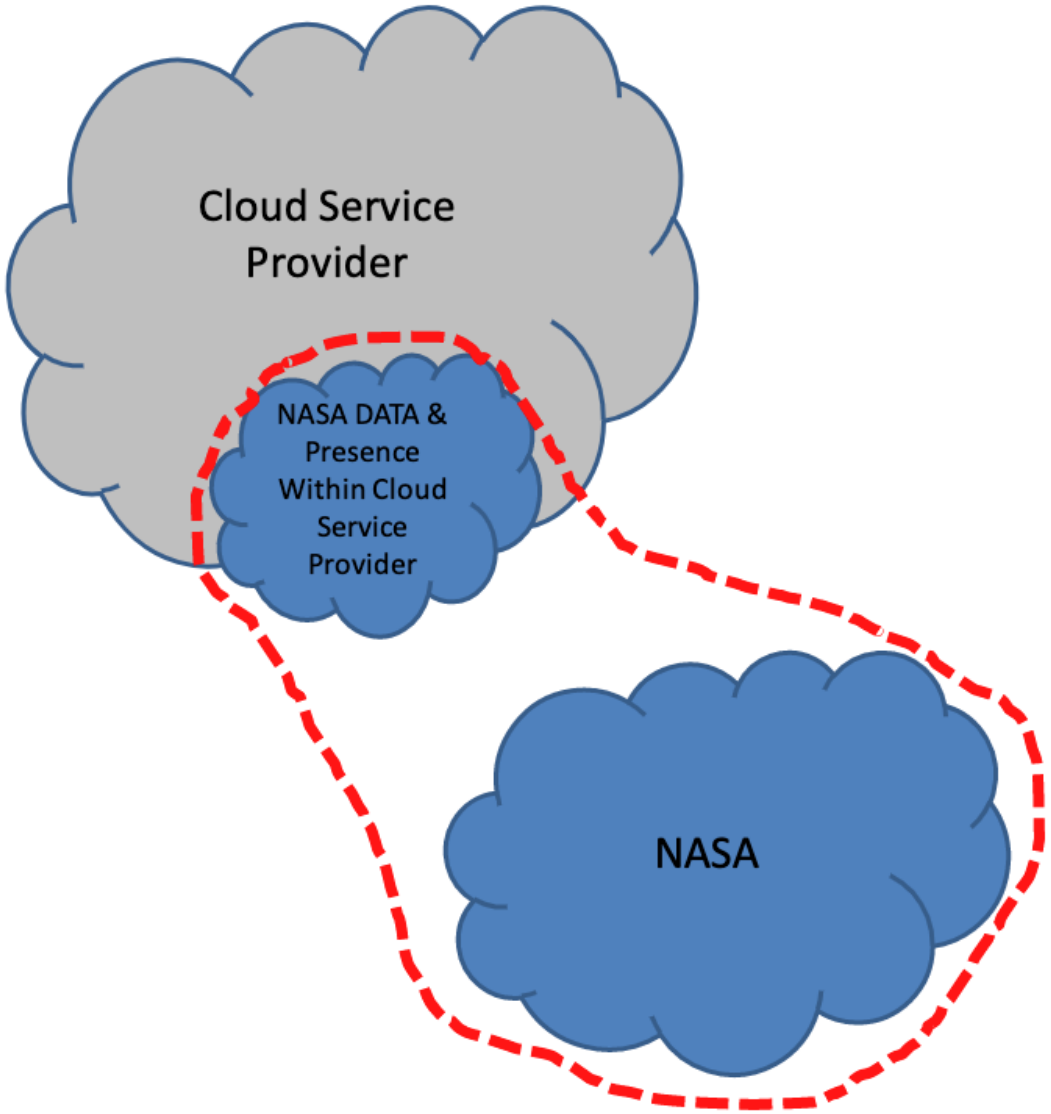

NASA Credit: https://ntrs.nasa.gov/citations/20170004723

What is Cloud Computing?

- What is Cloud Computing?
  - networks, servers, storage, applications, and services that you can rent
- Top Cloud Service Providers (CSP) in 2021
  - Amazon
  - Microsoft Azure
  - Alibaba Cloud
  - Google Cloud
  - Sales/force
  - Dell
  - IBM
  - Digital Ocean
  - Dropbox

What is AWS: Origins

- Evolved from Amazon's retail shop Jeff Bezos internal memo ("2002 Manifesto")

All teams will henceforth expose their data and functionality through service interfaces.

no other form of interprocess communication: no direct linking, no direct reads …, no shared-memory, no back-doors whatsoever….

…doesn't matter what technology ….all service interfaces must be externalizable to developers in the outside world ...

Anyone who doesn't do this will be fired.
have a nice day!

When should you consider cloud?

- When you have Requirements for:
  - Huge or Massive data
  - Fast real-time data processing
  - High traffic (Requests per Second (RPS))
  - Auto-failover even if physically destroyed (High Availability (HA))
  - Public access to NASA
    - Collaborate with Users or Industry outside of NASA
    - “non-NASA badged”
  - Avoid Procurement delays
    - Eg, Rent a GPU

Cloud Computing Pros and Cons

Pros	Cons
No procurement delays (once you are on board)	
Accessible from anywhere (subject to security)	NASA network plumbing can go down
Pay for what you use	Amazon manages infrastructure security and you manage application security (eg every resource is fire-walled)
Baked-in security and vulnerability patching	Vendor lock-in
Scalability - essentially unlimited	Ongoing costs can get out of control if not managed
Reliability – redundancy and backups are baked in	
Well-architected solutions to copy	Steep learning curve

Over 250 services! You don't know what you don't know

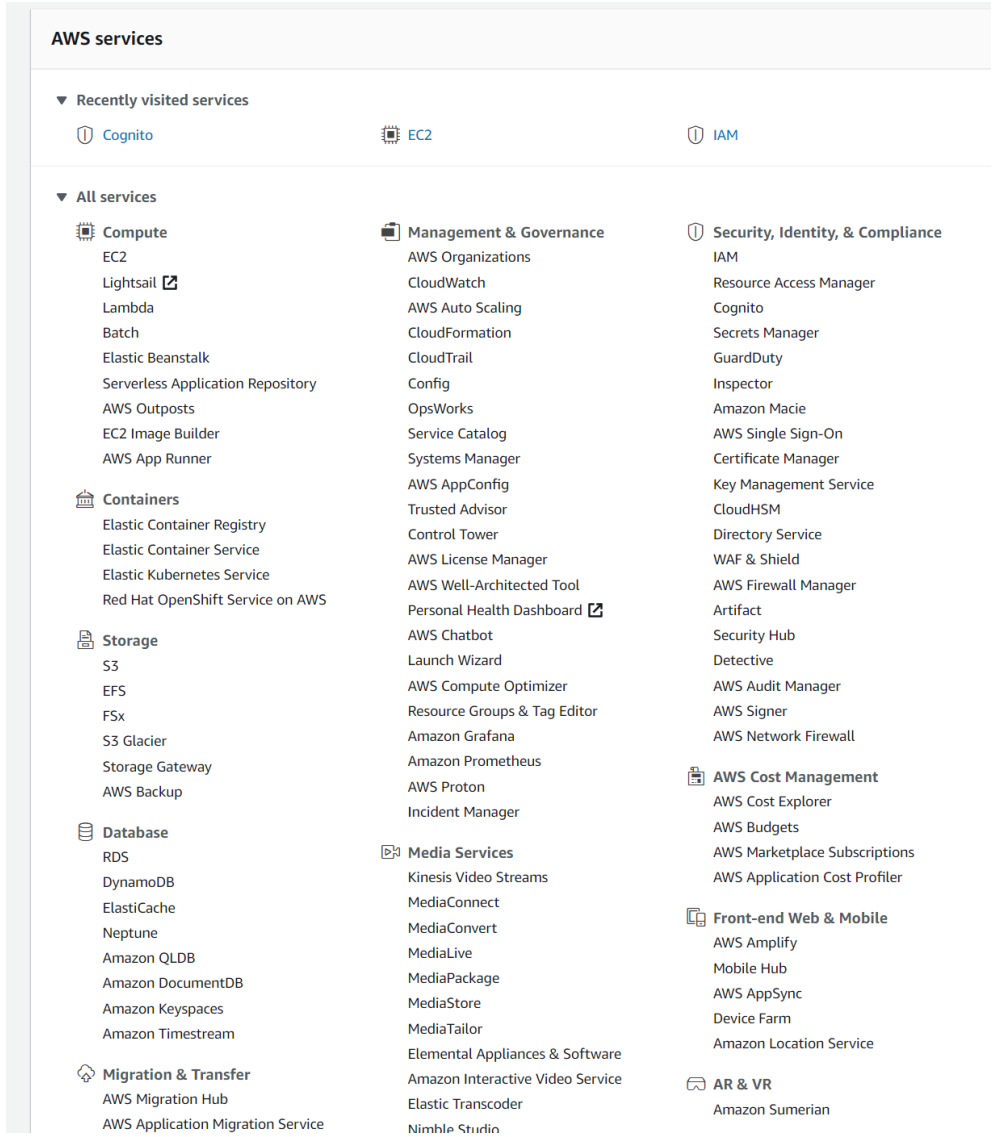

AWS services

Recently visited services

Cognito | EC2 | IAM

All services

Compute
EC2
Lightsail
Lambda
Batch
Elastic Beanstalk
Serverless Application Repository
AWS Outposts
EC2 Image Builder
AWS App Runner

Containers
Elastic Container Registry
Elastic Container Service
Elastic Kubernetes Service
Red Hat OpenShift Service on AWS

Storage
S3
EFS
FSx
S3 Glacier
Storage Gateway
AWS Backup

Database
RDS
DynamoDB
ElastiCache
Neptune
Amazon QLDB
Amazon DocumentDB
Amazon Keyspaces
Amazon Timestream

Migration & Transfer
AWS Migration Hub
AWS Application Migration Service

Management & Governance
AWS Organizations
CloudWatch
AWS Auto Scaling
CloudFormation
CloudTrail
Config
OpsWorks
Service Catalog
Systems Manager
AWS AppConfig
Trusted Advisor
Control Tower
AWS License Manager
AWS Well-Architected Tool
Personal Health Dashboard
AWS Chatbot
Launch Wizard
AWS Compute Optimizer
Resource Groups & Tag Editor
Amazon Grafana
Amazon Prometheus
AWS Proton
Incident Manager

Media Services
Kinesis Video Streams
MediaConnect
MediaConvert
MediaLive
MediaPackage
MediaStore
MediaTailor
Elemental Appliances & Software
Amazon Interactive Video Service
Elastic Transcoder
Nimble Studio

Security, Identity, & Compliance
IAM
Resource Access Manager
Cognito
Secrets Manager
GuardDuty
Inspector
Amazon Macie
AWS Single Sign-On
Certificate Manager
Key Management Service
CloudHSM
Directory Service
WAF & Shield
AWS Firewall Manager
Artifact
Security Hub
Detective
AWS Audit Manager
AWS Signer
AWS Network Firewall

AWS Cost Management
AWS Cost Explorer
AWS Budgets
AWS Marketplace Subscriptions
AWS Application Cost Profiler

Front-end Web & Mobile
AWS Amplify
Mobile Hub
AWS AppSync
Device Farm
Amazon Location Service

AR & VR
Amazon Sumerian

AWS Migration Hub
AWS Application Migration Service
Application Discovery Service
Database Migration Service
Server Migration Service
AWS Transfer Family
AWS Snow Family
DataSync

Networking & Content Delivery
VPC
CloudFront
Route 53
API Gateway
Direct Connect
AWS App Mesh
AWS Cloud Map
Global Accelerator

Developer Tools
CodeStar
CodeCommit
CodeArtifact
CodeBuild
CodeDeploy
CodePipeline
Cloud9
CloudShell
X-Ray
AWS FIS

Customer Enablement
AWS IQ
Support
Managed Services
Activate for Startups

Robotics
AWS RoboMaker

Blockchain
Amazon Managed Blockchain

Satellite
Ground Station

Quantum Technologies
Amazon Braket

Elastic Transcoder
Nimble Studio

Machine Learning
Amazon SageMaker
Amazon Augmented AI
Amazon CodeGuru
Amazon DevOps Guru
Amazon Comprehend
Amazon Forecast
Amazon Fraud Detector
Amazon Kendra
Amazon Lex
Amazon Personalize
Amazon Polly
Amazon Rekognition
Amazon Textract
Amazon Transcribe
Amazon Translate
AWS DeepComposer
AWS DeepLens
AWS DeepRacer
AWS Panorama
Amazon Monitron
Amazon HealthLake
Amazon Lookout for Vision
Amazon Lookout for Equipment
Amazon Lookout for Metrics

Analytics
Athena
Amazon Redshift
EMR
CloudSearch
Elasticsearch Service
Kinesis
QuickSight
Data Pipeline
AWS Data Exchange
AWS Glue
AWS Lake Formation
MSK
AWS Glue DataBrew
Amazon FinSpace

Amazon Sumerian

Application Integration
Step Functions
Amazon AppFlow
Amazon EventBridge
Amazon MQ
Simple Notification Service
Simple Queue Service
SWF
Managed Apache Airflow

Business Applications
Amazon Connect
Amazon Pinpoint
Amazon Honeycode
Amazon Chime
Amazon Simple Email Service
Amazon WorkDocs
Amazon WorkMail
Alexa for Business

End User Computing
WorkSpaces
AppStream 2.0
WorkLink

Internet of Things
IoT Core
FreeRTOS
IoT 1-Click
IoT Analytics
IoT Device Defender
IoT Device Management
IoT Events
IoT Greengrass
IoT SiteWise
IoT Things Graph

Game Development
Amazon GameLift

AWS Commonly Used Services

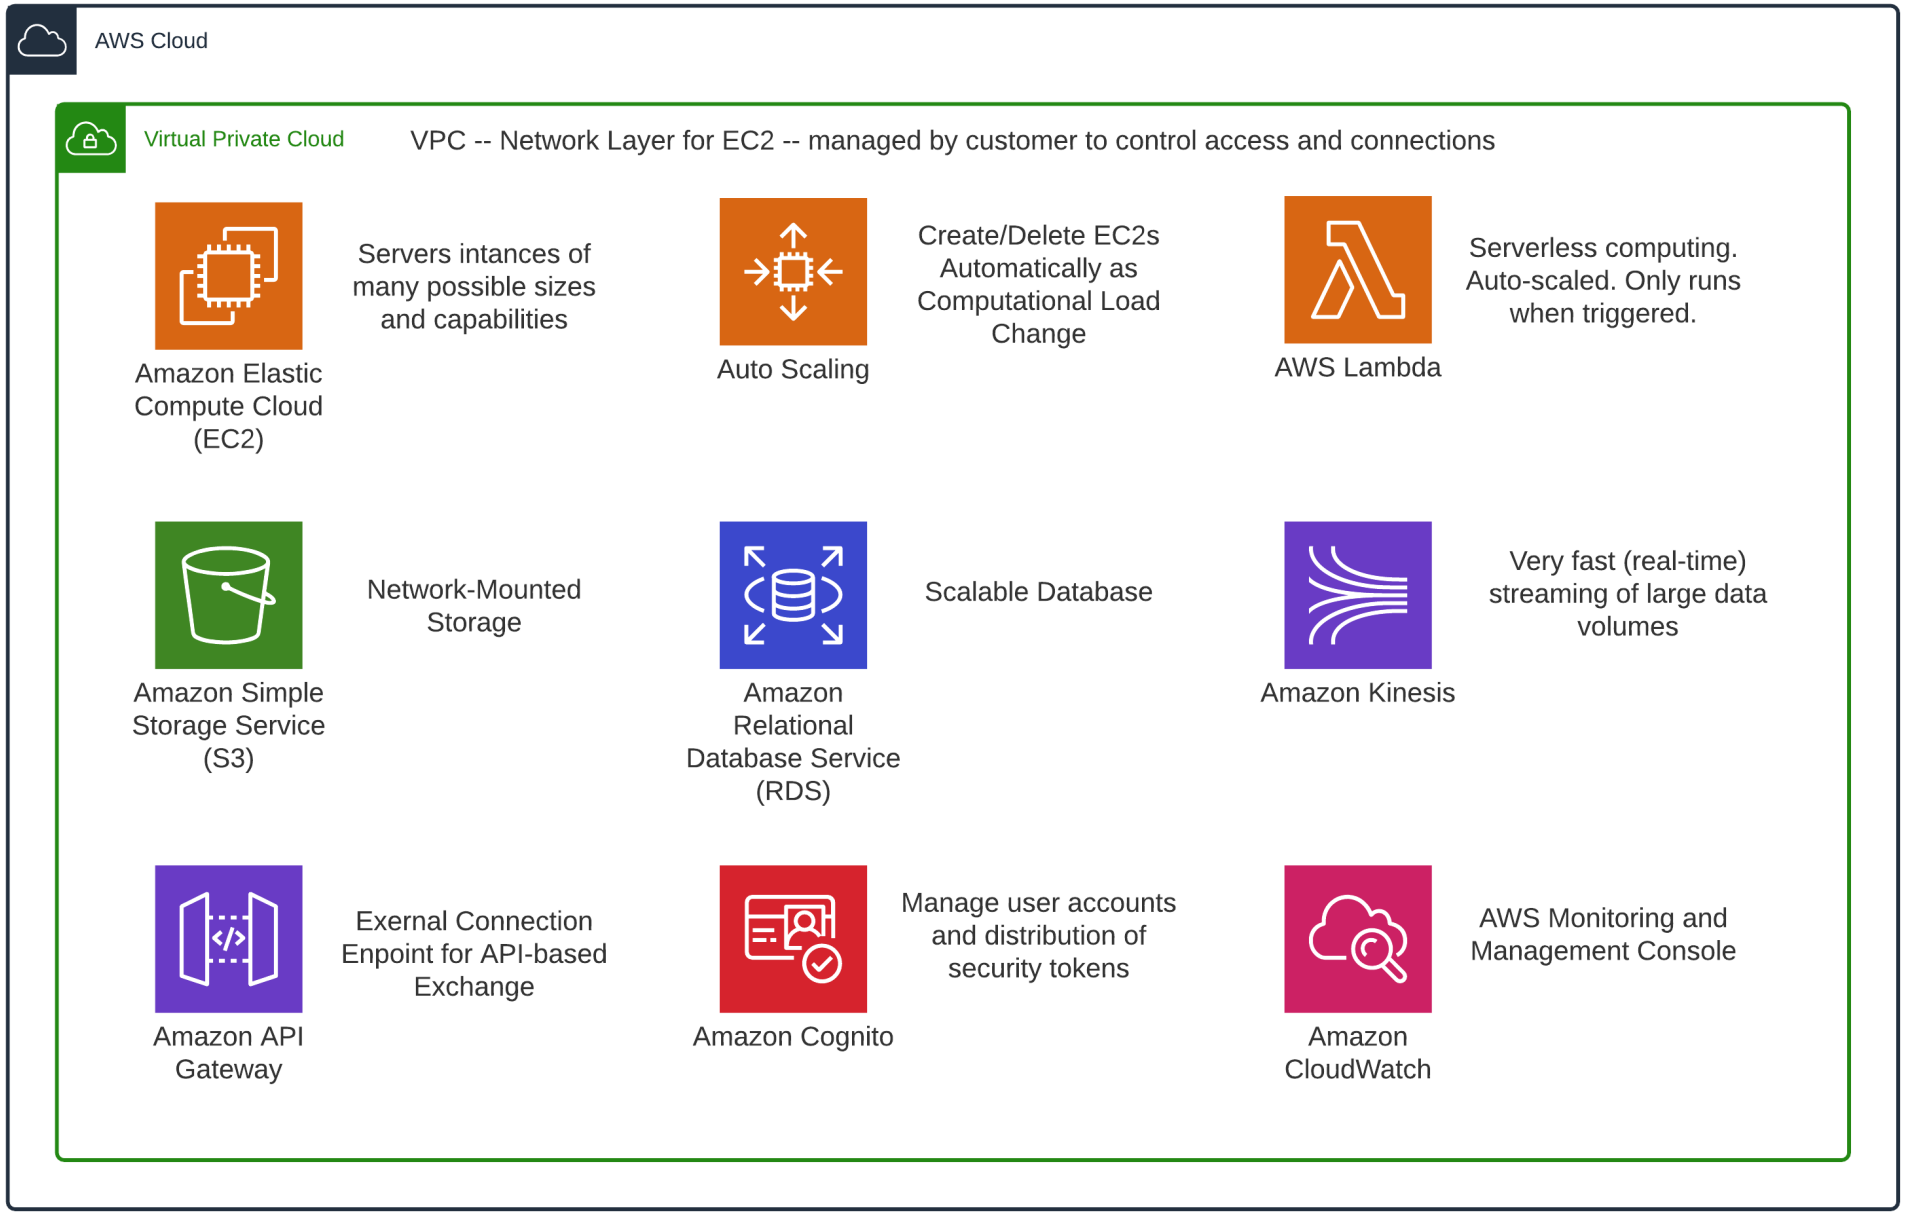

How Do We Manage Our AWS?

- Infrastructure as Code (IaC)
  - Errg…. Bob installed v1.2 on all machines, but later inadvertently Alice installed v2.05 on one of the machines! (Infrastructure Drift)
  - Provision infrastructure and applications through machine-readable configuration (text) files
  - Text files are SCM'ed
  - Deployments have audit log

Security in AWS

- Amazon provides security OF the cloud; customer provides security IN the cloud
- Zero Trust Model, aka "swiss cheese layers". Assume attacker is already in
- NASA IT Cloud Orgs approve all NASA AWS usage. You can’t just create AWS instances with a P-Card (anymore!)
- Three areas of Security:
  - Identity and Access Management (IAM) – users and their privileges.
  - Network Security – VPCs contain subnets; NASA-managed
  - Data Encryption

AWS Within NASA

- Looking back at NASA: 2009 Nebula, 2013 Cloud First, 2017 EMCC started
- NASA Enterprise Managed Cloud Computing (EMCC)
  - EMCC is a collection of organizations and teams across NASA that assure NASA policies are enforced, from security to finances
  - Code AF absorbed some of the “growing pains” as we started working with EMCC in 2018
  - Cloud at NASA is not all puppies and sunshine. But neither is running your own
  - Twice since 2018 NASA-plumbed accounts failed due to NASA network-down
  - In July 2021, Data Pipeline stopped working due to a security policy change by NICS at MSFC
  - Code AF has growing pains, too
- Code AF hired its own Cloud Engineer who is our Point of Contact to other orgs:
  - EMCC
  - NWMT: review/approve NASA Domain Name System (DNS) entries
  - AART: pen-tests, review applications before going public
  - Web Service Office: webapps, for example Fortigate Content Management
  - The Corporate Network Operations Center (CNOC) maintains Goddard Space Flight Center (GSFC) Trusted Internet Connection (TIC). The TIC supports VPN

AWS Within NASA

- EMCC provides the connections (“plumbing”) between NASA and AWS.
- Cloud vender agnostic
- NASA policy is that all Public sites need a domain name mapping ending in nasa.gov

NASA subdomain: nasa.gov

New!

Code AF: amesaero.nasa.gov

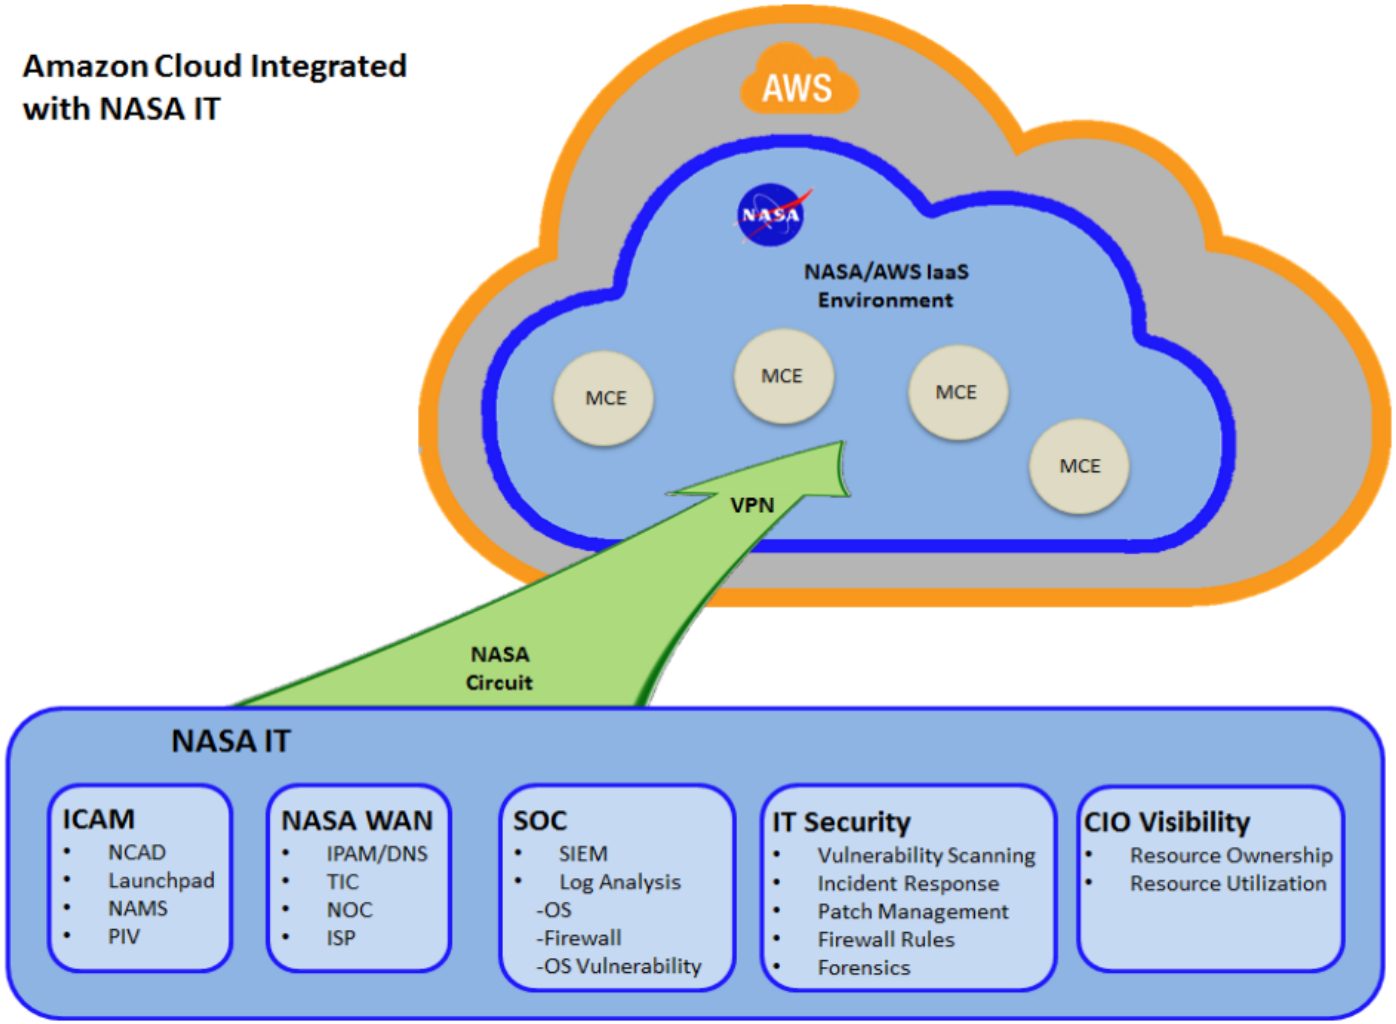

NASA Credit: https://ntrs.nasa.gov/citations/20170004723

Airspace Ops Management (AOM) Use of AWS

- AOM system runs a hybrid of AWS and NASA (on-premises) services that interact with each other
  - Data Pipeline and Real-Time Database in AWS
  - Grafana Visualization runs on NASA resources
  - Aerograph post-flight analysis runs on NASA resources, and data is stored on premises for long term, a common pattern
- ATM-X also uses a mix of NASA and AWS resources
  - ATM-X AWS resources are mostly EC2 compute servers
  - Remainder are Code AF servers or laptops
- Many security controls applied at interconnection points, especially if either side is ‘public’, that is, outside the NASA security boundary
- AOM Data Pipeline uses Serverless architecture

Data Pipeline's AWS Components – Real-Time System

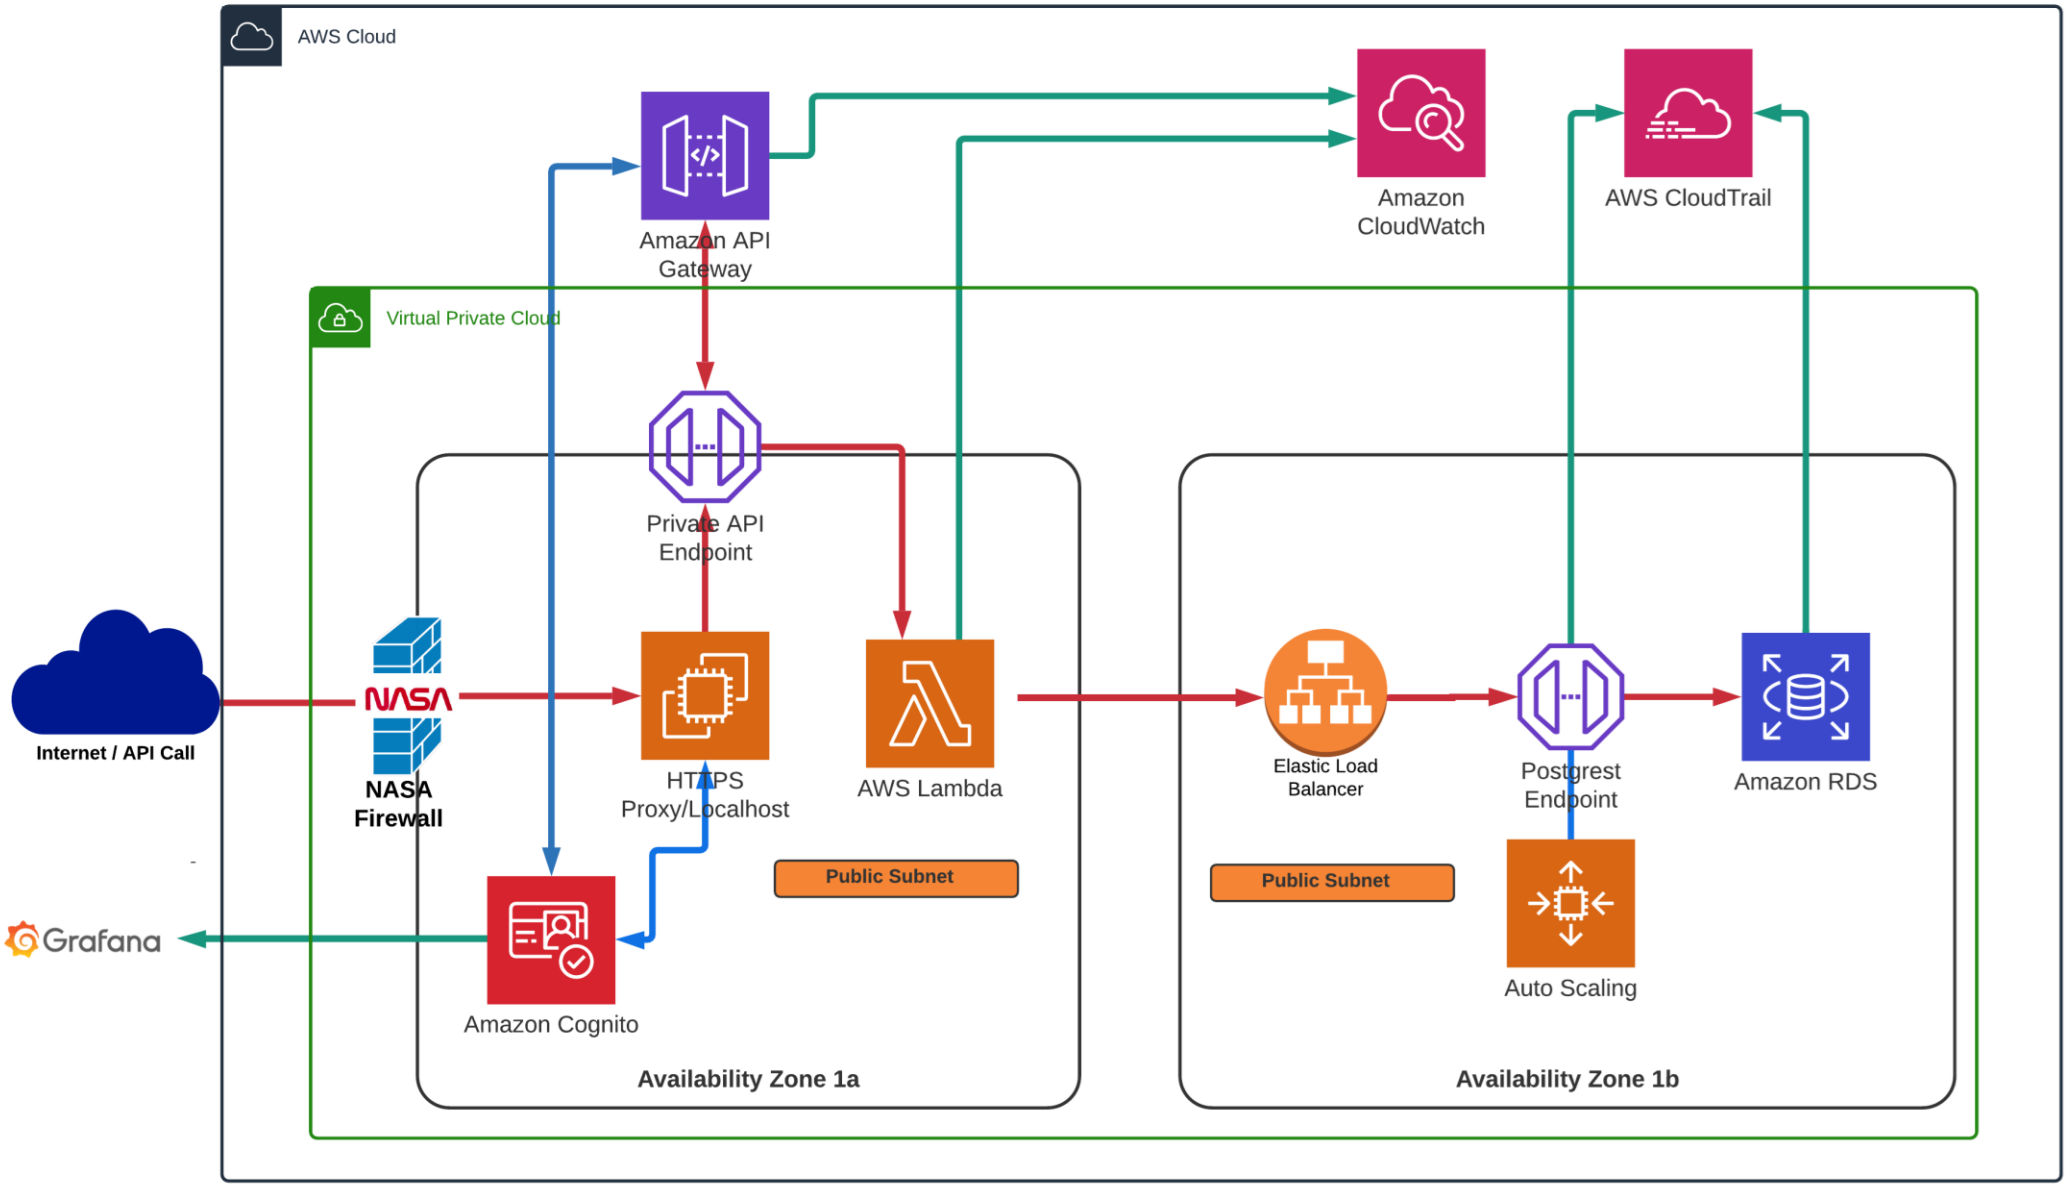

Serverless vs Traditional Cloud Virtual Machines

- **Compute:** Rather than running a virtual machine (EC2) for days or months, your code (lambda) runs in a pool of EC2's whose lifetime is roughly the duration is the run-time of your code. (e.g. 3 seconds)
- **Cost-efficiency:** 'pay-as-you-go' plans. For EC2s, you pay by the hour no matter if it is used or not. For Serverless compute you for pay per Request, or per milli-seconds of use
- **Maintenance:** no need to patch Operating System. Check at design time that your service is approved for federal government use (FedRamp).
- Rather than setting up for High Availability and Scaling yourself, these concerns are baked in

AOM Continuous Provisioning and Deployment

- Developer peer reviews code changes then commits to ‘git’ SCM
- Jenkins triggers can be per-commit or scheduled; We run system tests nightly to catch regressions
- AWS CloudFormation provisions infrastructure and the serverless application based on parameters Environment (dev/stage/production), and an EventID discriminator to Stratify data by Owner
  - Infrastructure is provisioned
  - Application is deployed
  - System tests are run nightly

Questions

APPENDIX

Cloud Computing NIST Definition

- What is Cloud Computing?
  - From US NIST:
    - Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model is composed of five essential characteristics, three service models, and four deployment models:

Essential Characteristics	Service Models	Deployment Models
On-demand self-service	Software as a Service (SaaS)	Private cloud
Broad Network Access	Platform as a Service (PaaS)	Community cloud
Resource Pooling	Infrastructure as a Service (IaaS)	Public cloud
Rapid Elasticity		Hybrid cloud
Measured Service		

AWS meets all NIST characteristics, supports PaaS and IaaS, and is a Public cloud (including gov cloud)

AWS Service Areas

- Compute
- Containers
- Storage
- Database
- Migration and Transfer
- Networking and Content Delivery
- Developer Tools
- Customer Enablement
- Robotics
- Blockchain
- Satellite
- Quantum Technologies
- Management and Governance
- Database
- Machine Learning
- Analytics
- Developer Tools
- Security, Identity & Compliance
- AWS Cost Management
- Front-end Web & Mobile
- AR & VR
- Application Integration
- Business Applications
- End User Computing
- Internet of Things
- Game Development

Security in AWS

- Amazon is responsible for security OF the cloud
- Customer is responsible for security IN the cloud
- A Zero Trust model is adopted – all application components and services are considered discrete and potentially malicious
- NASA IT Cloud Orgs control, manage, and approve all NASA AWS usage. A project can't just create AWS instances with a P-Card (anymore!)
- Three areas of Security:
  - Identity and Access Management (IAM) – track identities, grant access – highly controlled by NASA
  - Network Security – system, configurations, and processes to safeguard access and usability of network and network-accessible resources. Lowest level component is the VPC – Virtual Private Cloud. Completely controlled by NASA Cloud Orgs
  - Data Encryption – at rest, in transit, as desired by project

References

EMCC: https://ntrs.nasa.gov/citations/20170004723

2013 NASA “Cloud First” initiative, FedRAMP, Nebula in 2009:
https://oig.nasa.gov/docs/IG-13-021.pdf