UCSF School of Medicine

Accelerating Secure Medical Research in the Cloud

Jamie Lam, Data Security Compliance Manager
Pavan Gupta, Cloud Architect
Victor Vargas Reyes, Software Engineering Manager

Agenda

1. Background and Compliance Approach
2. Architecture, Tools
3. Operations
4. Looking Ahead

Who is UCSF SOM Tech?

We are technology advocates. Our broad technology expertise, deep knowledge of the UCSF ecosystem, and human-centered approach help you take innovative digital projects further, faster.

UCSF School of Medicine

Background and Compliance Approach

How did this all start?

Demand

- Rising demand for cloud computing **with PHI** (multiple inquiries / month and increasing)

Cost

- >$500k spend at UCSF / year on AWS

Security

- No visibility into existing accounts
- No controls vetting researchers using cloud infrastructure

Governance Accelerated Path Forward

Enterprise Hybrid Cloud	Departmentally Managed Public Cloud
Hybrid cloud hosting is a model that allows UCSF to provision dedicated servers and storage and shared cloud servers and storage on the same network. UCSF IT will offer a cloud service catalog.	In public clouds, cloud resources are owned and operated by a third-party cloud service provider and delivered over the Internet. Departments are responsible for managing their own public clouds.
• Business requirement to connect to UCSF network / systems, and/or • Department wants to leverage enterprise security tools that will be integrated to the enterprise hybrid cloud infrastructure	• No business requirement to have persistent connectivity to UCSF network / systems, and • Department has resources and expertise to configure and operationally manage a secure public cloud infrastructure
• Identify and commit to timeline, resource need, and tools (such as DirectConnect) required • Funding for resources and tools • Define service lines available • Define roles and responsibilities for UCSF IT vs. Department IT under this model • Other requirements may be defined as we learn	• Department must document architecture design and correlate how cloud tools and configurations are used to meet HIPAA/CSA security controls • Department must draft IT Security Plan with policies and procedures to operationally manage the public cloud • Other requirements may be defined as we learn
UCSF can better leverage existing enterprise security tools and will have more visibility into our cloud infrastructure. Enterprise cloud service line allows departments and PIs without cloud security expertise to consume cloud services securely.	Can get started now and provides department an option to roll their own. Multiple deployments of public clouds will result in inefficiency and added resource requirements with each department deploying their own SIEM/ host based detection/ VPN services, etc.

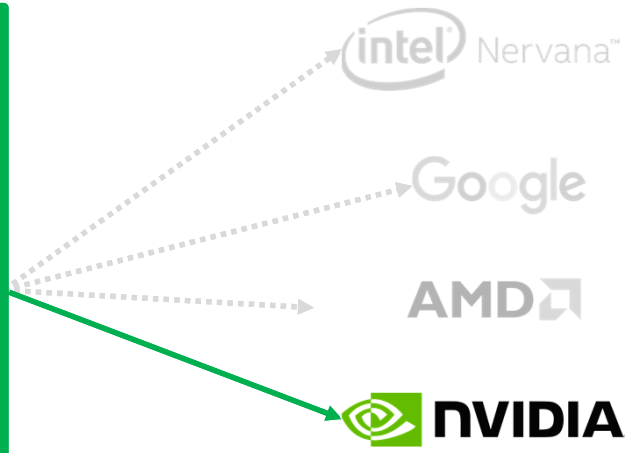

Cloud Computing Provides:

- **Hardware Access**
- **Hardware Optionality**
- Compute Flexibility

Compelling Research Use Cases

Four Early Use Cases:

1. Linux Research Computing (FAST and Pregnancy Ultrasound Research)
2. Containerized Computing (UCSB's BisQue Platform, HPC Analog)
3. De-identified clinical data warehouse (Information Commons)
4. Specialized Clinical Applications (CDHI's Fax-2-Referral)

Initial Team

- SOM Tech Data Security Compliance Manager
- SOM Tech Project Manager
- CDHI AWS Architect
- ClearScale Architect
- ClearScale DevOps Engineers
- UCSF Enterprise:
  - IT Security
  - Privacy

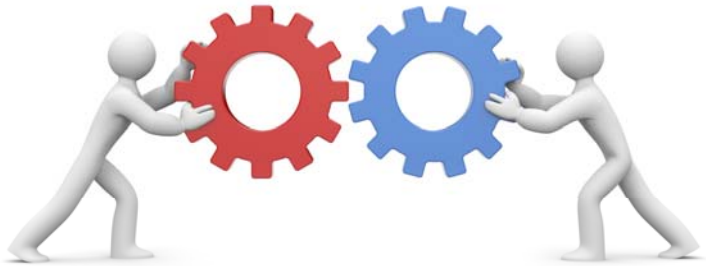

Compliance

- Cloud Security Alliance Cloud Controls Matrix
  - ~130 cloud specific security controls mapped to industry security standards / regulations such as ISO 27001/27002, NIST, and HIPAA
  - Mapped each control to an operational procedure
  - Incorporated as technical design requirements
- UCOP IS-3 Electronic Information Security
  - Drafted IT Security Plan for environment
- UCSF Policy 650-16 (Minimum Security Standards)
  - Standards pushed into technical design
- UCOP-AWS Business Associate Agreement (BAA)

Compliance Documentation

- IT Security Plan
- 12 procedures

1. Risk Management
2. Information Access Management
3. Systems Inventory, Data Management & Retention
4. Configuration Management
5. Encryption & Key Management
6. Vulnerability Scanning & Management
7. Change Management
8. Logging and Monitoring
9. Business Continuity & Disaster Recovery
10. Incident Response
11. Security Awareness and Training
12. Physical Security

Sprinting Toward Actual Operations

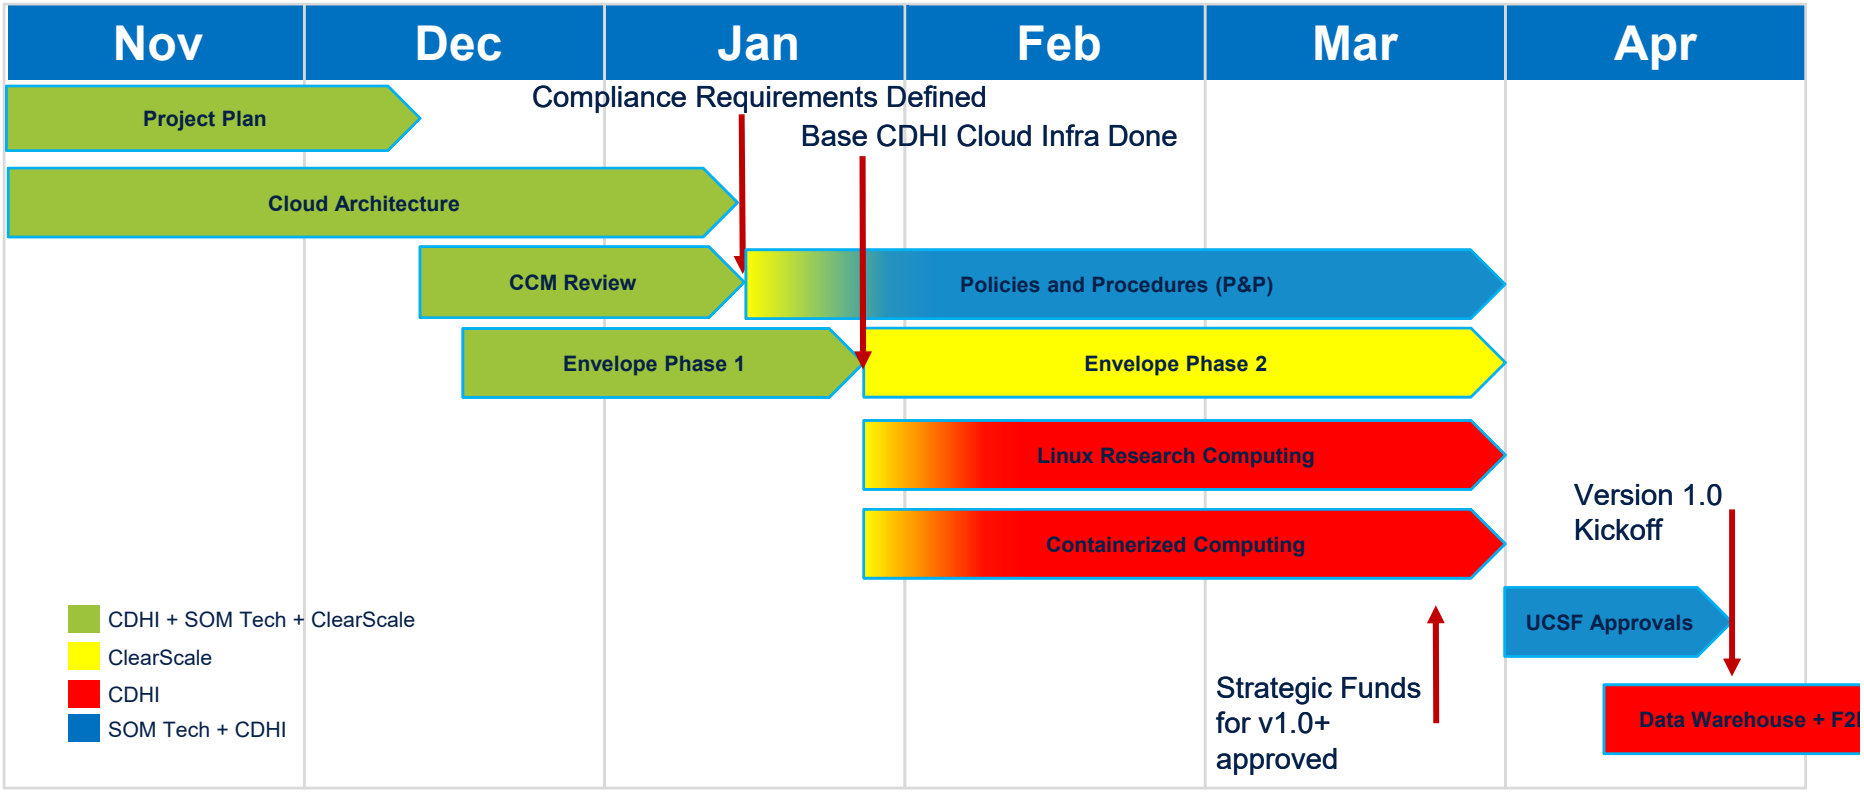

UCSF School of Medicine

Architecture

Architecture

Security

Operations

Projects

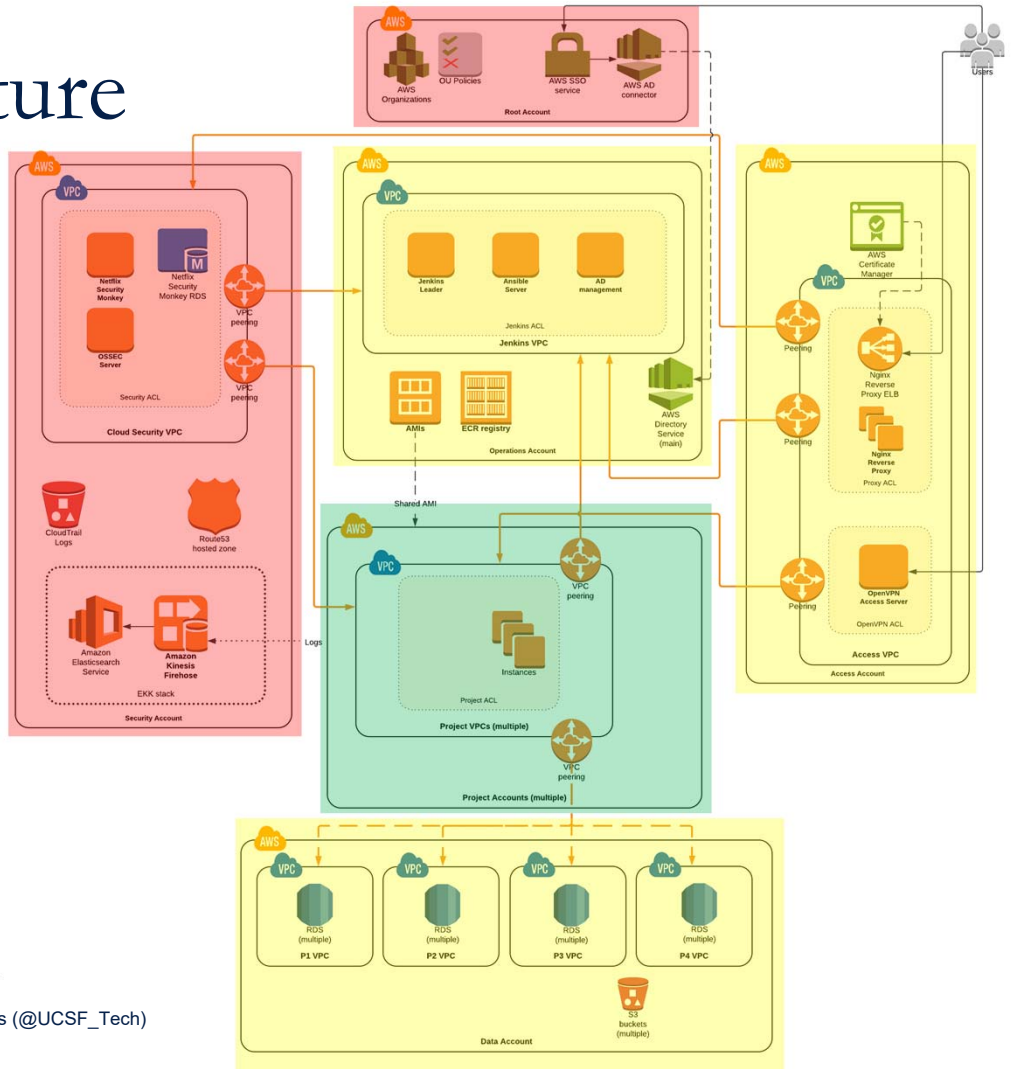

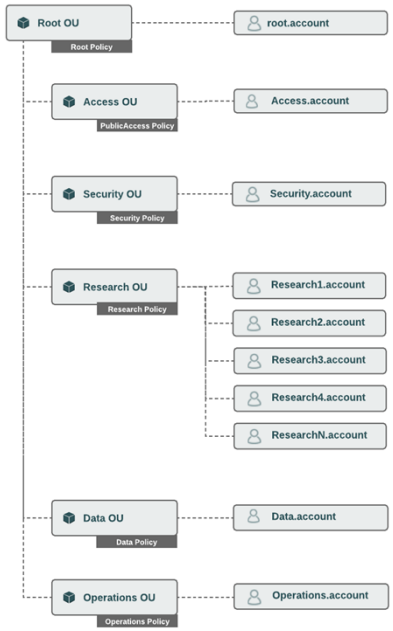

VPN and Web Connectivity

Goal:

Single point of entry, RBAC, auditable and extremely secure, layers of additional web-related security

Current Challenges:

OpenVPN is not the enterprise choice

Security (host posture checking, dlp) not used

Route53 used with cdhi.cloud right now

Nginx Reverse Proxy != MyAccess SSO

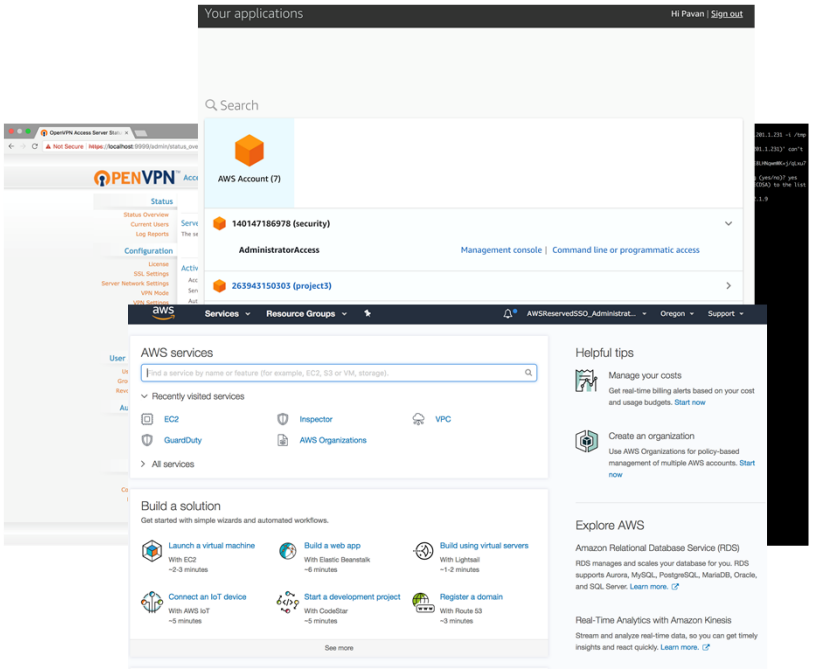

Compliance/Security Tooling

Goal:

Sophisticated tools to ensure auditability and active monitoring of cloud deployment

Current Challenges:

Tools lack enterprise bells and whistles

Alerting is not effectively enabled

Better tools may exist within UCSF

Dome9 is still experimental within our platform

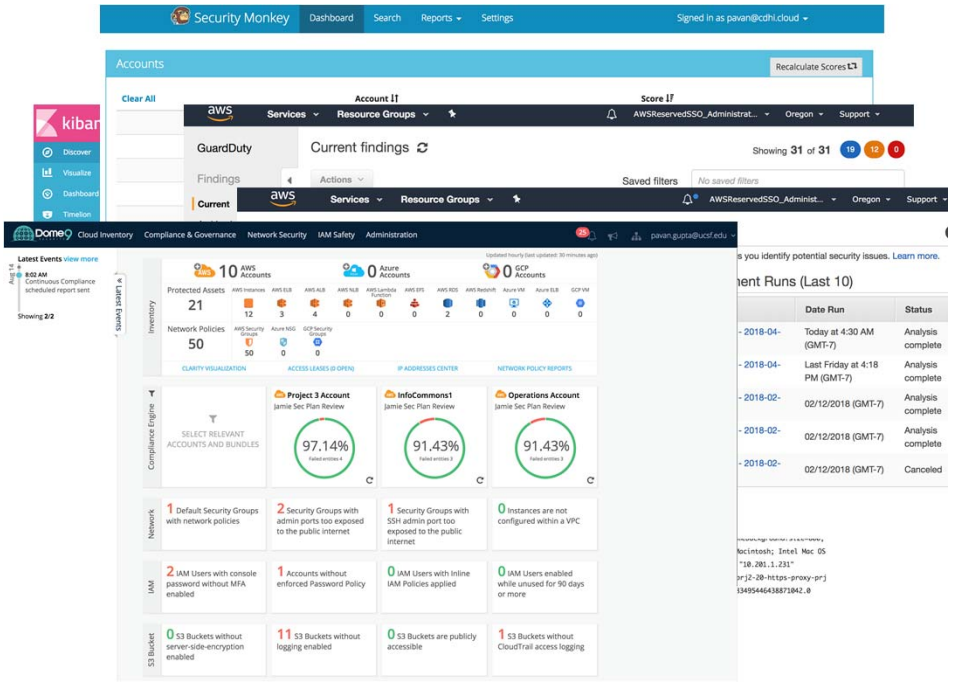

Data Management

Goals:

Extremely segmented, manual distribution of secure data

Current Challenges:

Integration with enterprise data stores not used

Data availability not guaranteed

Vanilla data stores being used

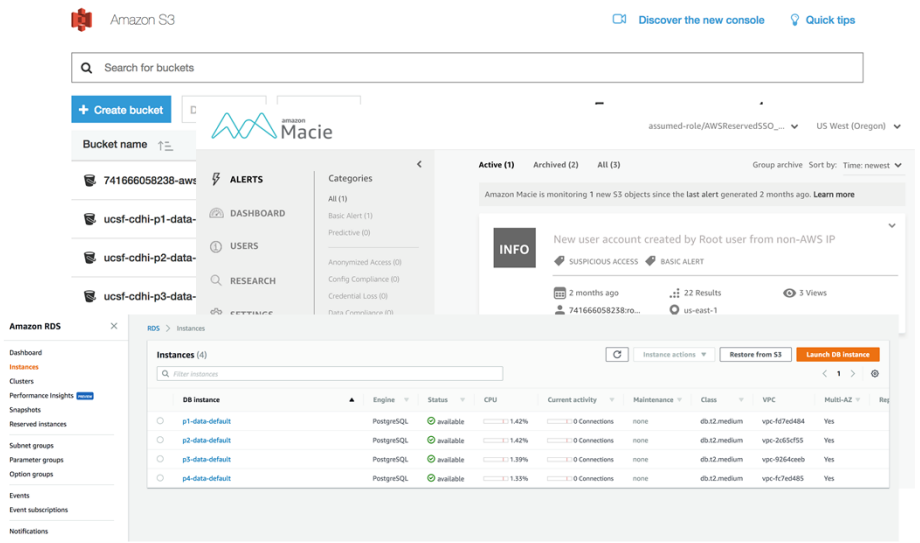

Account, User, & Billing Management

Goal:

Centrally managed, centrally organized account and billing structure requiring minimal direct maintenance

Current Challenges:

IAM + Organizations Policies can be improved

No multi-factor authentication available

Another set of credentials required for users

Manual validation of all active users required

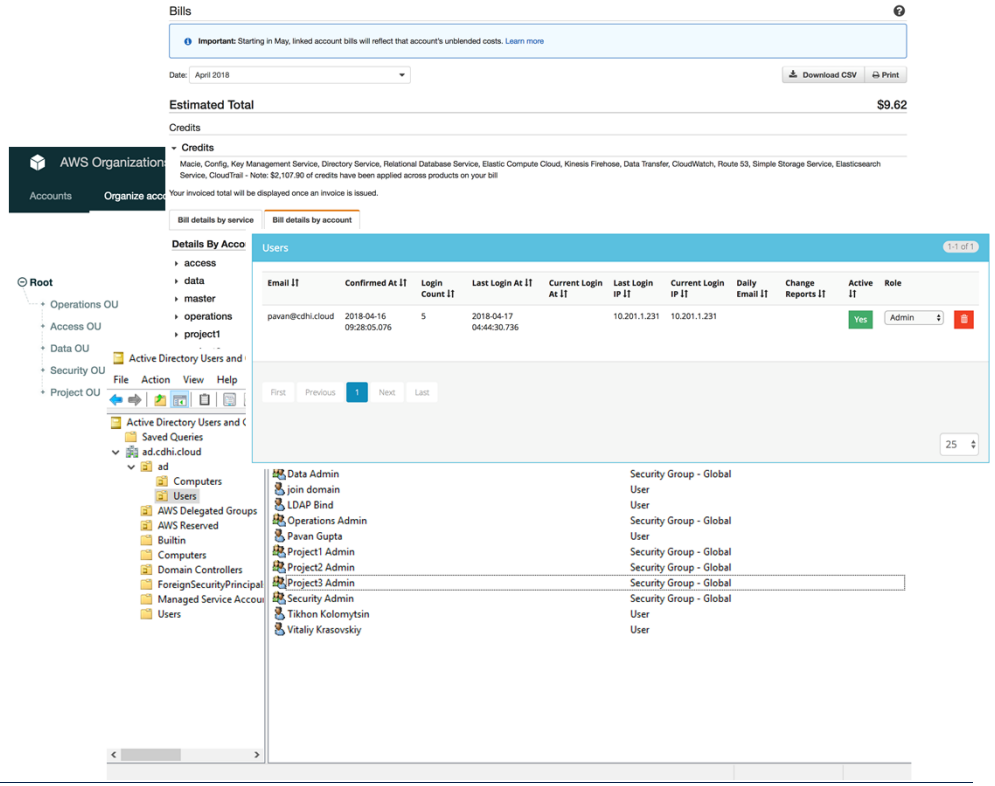

Operations/Change/Config Management

Goals:

Centrally managed, integrated change and config management strategy

Current Challenges:

ServiceNow integration not available

JIRA + Confluence not the best CM option

Only one baseline image available

Match enterprise configuration management tools

Terraform/Packer/Docker/Ansible are not simple

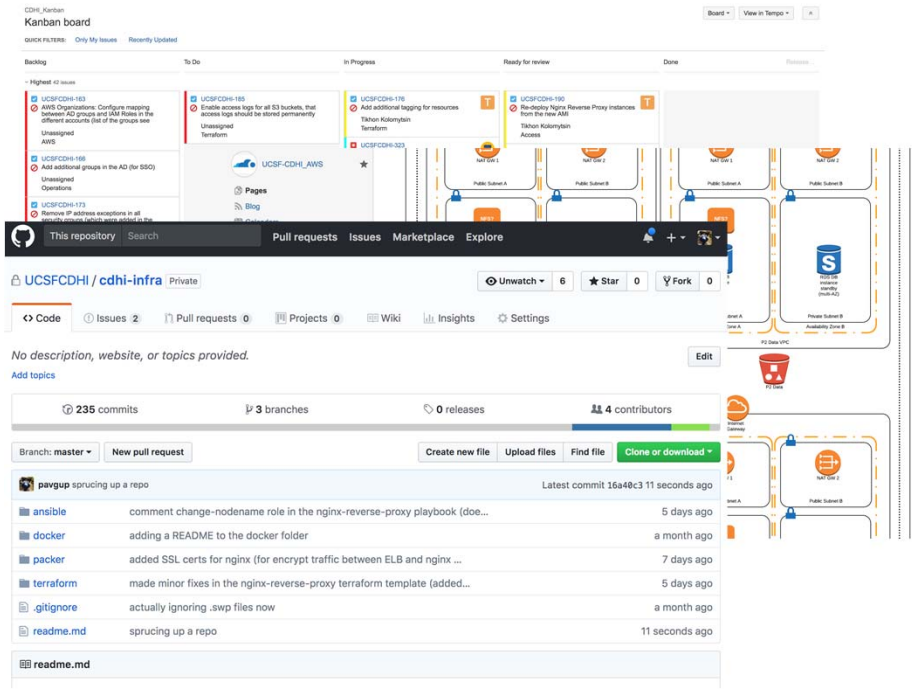

Real User Experience

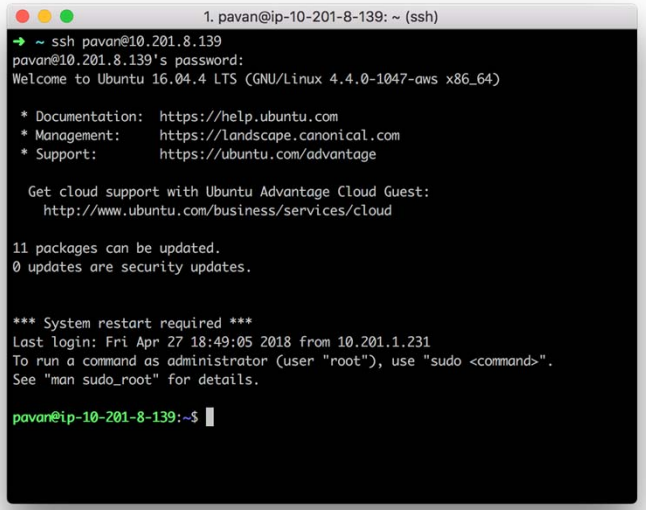

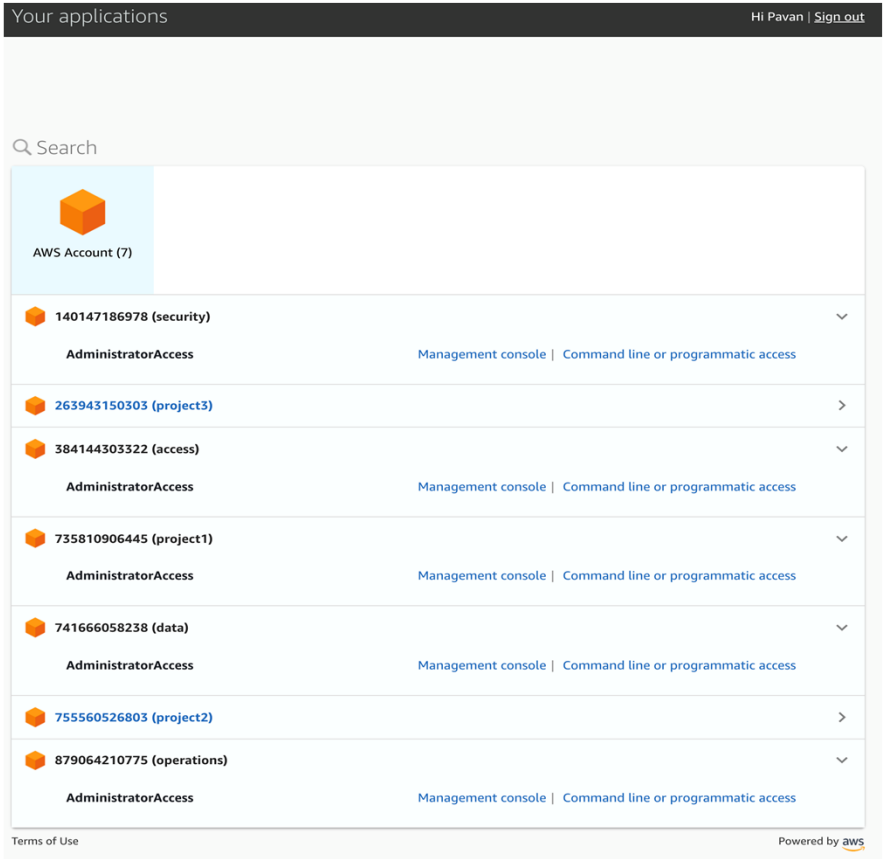

Sprinting Toward Actual Operations

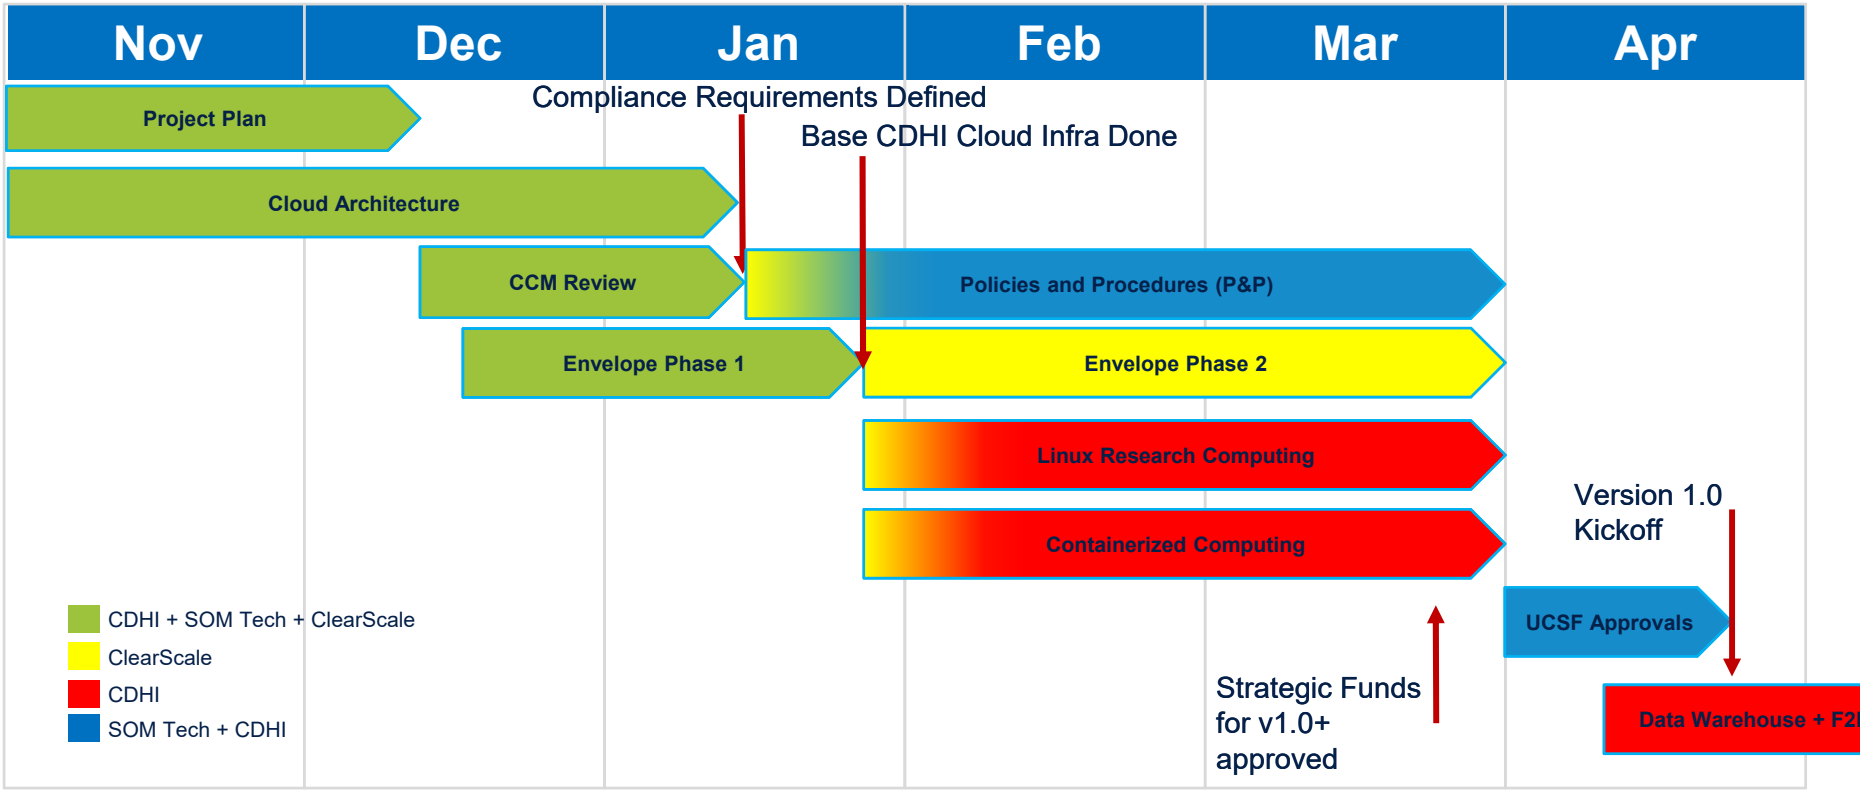

UCSF School of Medicine

Operations

Ops, Ops, Ops and More Ops

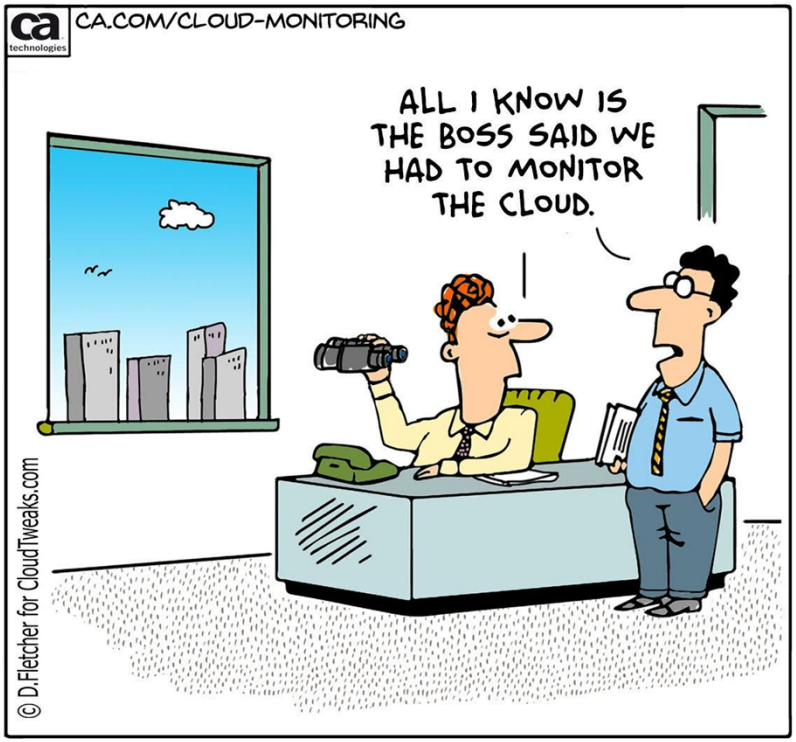

Operations – Technology

- Terraform
- Packer
- Ansible
- GitHub
- Jenkins Pipelines
- ServiceNow
- Jira

Operations – Procedures and Support

- Compliance Tasks
  - Dome9 Review
  - AD Review
  - Network Review
  - Privileged account use
  - Inspector Reviews
  - Encryption
  - And more…
- Operations Tasks
  - AMI Builds
  - Resource Tagging
  - Remediation
  - Provisioning
  - IaC Tech Debt
  - Billing
  - And more…

Operations – Tech Debt Management

- Agile Project
  - Migrate infrastructure
  - New features
  - New automation pieces
- DevOps Team
  - .5 FTE for Operations
  - 1.5 FTE for Tech Debt and Remediation

Operations – Procedures and Team

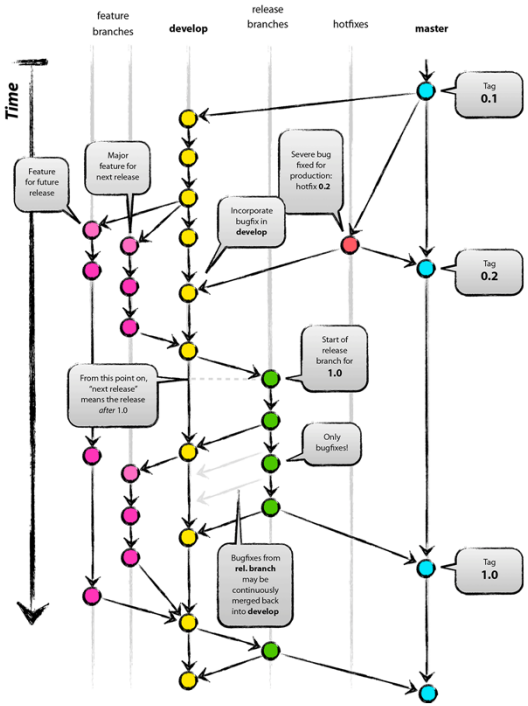

- Change Management
  - Testing
  - Code Review
  - Peer Review
- SCM
  - Pull Requests
  - Protected Branches
  - Git flow

UCSF

Operations – Challenges for SOM Tech

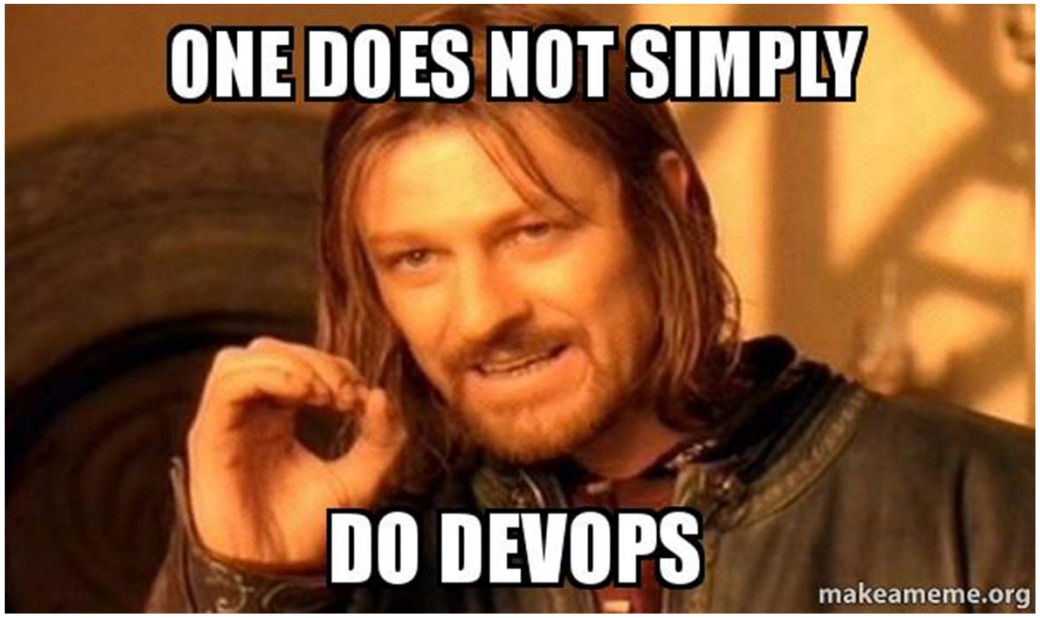

Operations – Challenges for SOM Tech

- New technologies with steep learning curve
- New development life cycle
- Automation
- IaC vs direct configuration
- Scalability and Enterprise IT
- Handoff

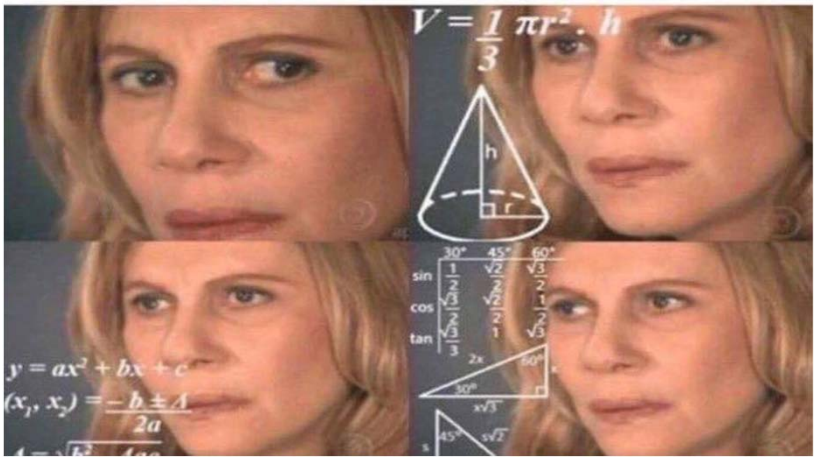

UCSF School of Medicine

Looking Forward

Strategically Looking Forward

- **Scalability**
  1. Operations and Security efficiency
  2. Enterprise integrations
  3. Billing
  4. Training

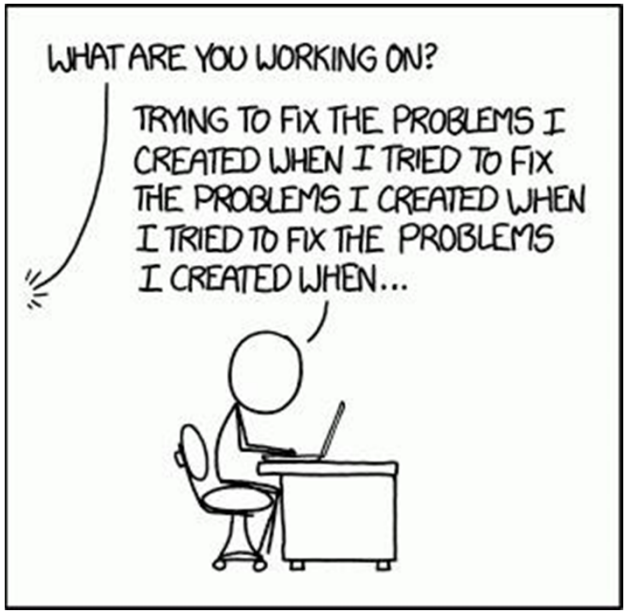

Strategically Looking Forward

- **Enabling Research**
  1. 30+ projects waiting for access, 90% research related
  2. Containerized computing with
  3. Serverless computing
  4. Advice and consulting

Acknowledgements

- **UCSF Enterprise:**
  - Pat Phelan, Chief Information Security Officer
  - Tom Poon, Interim Chief Privacy Officer
  - Toby Barber, Information Security Architect
  - Bianca Paraguya, Privacy Office
- **SOM Tech:**
  - Kristin Chu, Director SOM Tech
  - Kymberly Ainsworth, Business Process Manager
  - Sean Thomas, Project Manager
- **CDHI:**
  - Rachael Callcut, Associate Professor of Surgery, Director of Data Science
  - Ed Martin, Director of Technology
  - Joe Hesse, Director of Innovation