



CDL Infrastructure Migration to the AWS Cloud – 1 Year Later

Presented by:

Kurt Ewoldsen

Director of Infrastructure & Application
Support

The California Digital Library (CDL)

- **Mission:** The California Digital Library exists to support the University of California community's pursuit of scholarship and to extend the University's public service mission.
- **Vision:** The California Digital Library's vision is to elevate the digital library for UC so that it becomes "expansively global and deeply local". CDL will advance the digital transition of scholarly information in three spheres:
 - Access: Scholars will have access to the highest quality research collections worldwide through services that support and enable new scholarship and make it as open as possible.
 - Formats: CDL will support all digital formats throughout their life cycle with a full range of services, especially to surface UC's unique digital assets and collections.
 - Scale: Through partnerships and alliances, CDL will elevate services to the network level for maximum impact.

See more at <http://www.cdlib.org>

Mission: Accomplished



Physical Sun/Solaris

1/1/2010-7/1/2015

VMware

3/1/2010-6/1/2016

AWS

1/1/2014-7/1/2016

AWS migration proposed

9/11/2013

VMware environment build

3/1/2010

AWS account created

1/1/2014

Solaris shut down

7/1/2015

VMware migration start

10/1/2010

New IT manager

7/8/2012

AWS migration start

7/1/2014

VMware shut down

5/17/2016

2010

2011

2012

2013

2014

2015

2016

Year to Year Comparison

	UCCSC 2015	UCCSC 2016
Amazon EC2 Instances	~100	~180
Amazon RDS Instances	~50	~60
Amazon S3 Buckets	~40/~15TB	~40 / ~20TB
Amazon Accounts	2	4
Amazon Costs	\$20,000/month	\$30,000/month

Project Goals

- Improve infrastructure stability
- Avoid replacing aging HW/SW
- Reduce range of required technical knowledge
- Eliminate SUSE release schedule issues
- Consolidate environments on a single infrastructure
- Enable cost effective HA/DR capabilities
- Provide responsive, flexible environment
- Track & report individual service infrastructure costs
- Increase IAS resources/bandwidth

Project Goals

- ✓ Improve infrastructure stability
- ✓ Avoid replacing aging HW/SW
- ✓ Reduce range of required technical knowledge
- ✓ Eliminate SUSE release schedule issues
- ✓ Consolidate environments on a single infrastructure
- ✓ Enable cost effective HA/DR capabilities
- ✓ Provide responsive, flexible environment
- ✓ Track & report individual service infrastructure costs
- ✓ Increase IAS resources/bandwidth

Project Goals

- ✓ Improve infrastructure stability
- ✓ Avoid replacing aging HW/SW
- ✓ Reduce range of required technical knowledge
- ✓ Eliminate SUSE release schedule issues
- ✓ Consolidate environments on a single infrastructure
- ✓ Enable cost effective HA/DR capabilities
- ✓ Provide responsive, flexible environment
- ✓ Track & report individual service infrastructure costs
- ✓ Increase IAS resources/bandwidth

Project Goals

- ✓ Improve infrastructure stability
- ✓ Avoid replacing aging HW/SW
- ✓ Reduce range of required technical knowledge
- ✓ Eliminate SUSE release schedule issues
- ✓ Consolidate environments on a single infrastructure
- ✓ Enable cost effective HA/DR capabilities
- ✓ Provide responsive, flexible environment
- ✓ Track & report individual service infrastructure costs
- ✓ Increase IAS resources/bandwidth

Project Goals

- ✓ Improve infrastructure stability
- ✓ Avoid replacing aging HW/SW
- ✓ Reduce range of required technical knowledge
- ✓ Eliminate SUSE release schedule issues
- ✓ Consolidate environments on a single infrastructure
- ✓ Enable cost effective HA/DR capabilities
- ✓ Provide responsive, flexible environment
- ✓ Track & report individual service infrastructure costs
- ✓ Increase IAS resources/bandwidth

Project Goals

- ✓ Improve infrastructure stability
- ✓ Avoid replacing aging HW/SW
- ✓ Reduce range of required technical knowledge
- ✓ Eliminate SUSE release schedule issues
- ✓ Consolidate environments on a single infrastructure
- ✓ Enable cost effective HA/DR capabilities
- ✓ Provide responsive, flexible environment
- ✓ Track & report individual service infrastructure costs
- ✓ Increase IAS resources/bandwidth

Project Goals

- ✓ Improve infrastructure stability
- ✓ Avoid replacing aging HW/SW
- ✓ Reduce range of required technical knowledge
- ✓ Eliminate SUSE release schedule issues
- ✓ Consolidate environments on a single infrastructure
- ✓ Enable cost effective HA/DR capabilities
- ✓ Provide responsive, flexible environment
- ✓ Track & report individual service infrastructure costs
- ✓ Increase IAS resources/bandwidth

Project Goals

- ✓ Improve infrastructure stability
- ✓ Avoid replacing aging HW/SW
- ✓ Reduce range of required technical knowledge
- ✓ Eliminate SUSE release schedule issues
- ✓ Consolidate environments on a single infrastructure
- ✓ Enable cost effective DR capabilities
- ✓ Provide responsive, flexible environment
- ✓ Track & report individual service infrastructure costs
- ✓ Increase IAS resources/bandwidth

Project Goals

- ✓ Improve infrastructure stability
- ✓ Avoid replacing aging HW/SW
- ✓ Reduce range of required technical knowledge
- ✓ Eliminate SUSE release schedule issues
- ✓ Consolidate environments on a single infrastructure
- ✓ Enable cost effective DR capabilities
- ✓ Provide responsive, flexible environment
- ✓ Track & report individual service infrastructure costs
- ✓ Increase IAS resources/bandwidth


Project Goals

- ✓ Improve infrastructure stability
- ✓ Avoid replacing aging HW/SW
- ✓ Reduce range of required technical knowledge
- ✓ Eliminate SUSE release schedule issues
- ✓ Consolidate environments on a single infrastructure
- ✓ Enable cost effective DR capabilities
- ✓ Provide responsive, flexible environment
- ✓ Track & report individual service infrastructure costs
- ✓ Increase IAS resources/bandwidth

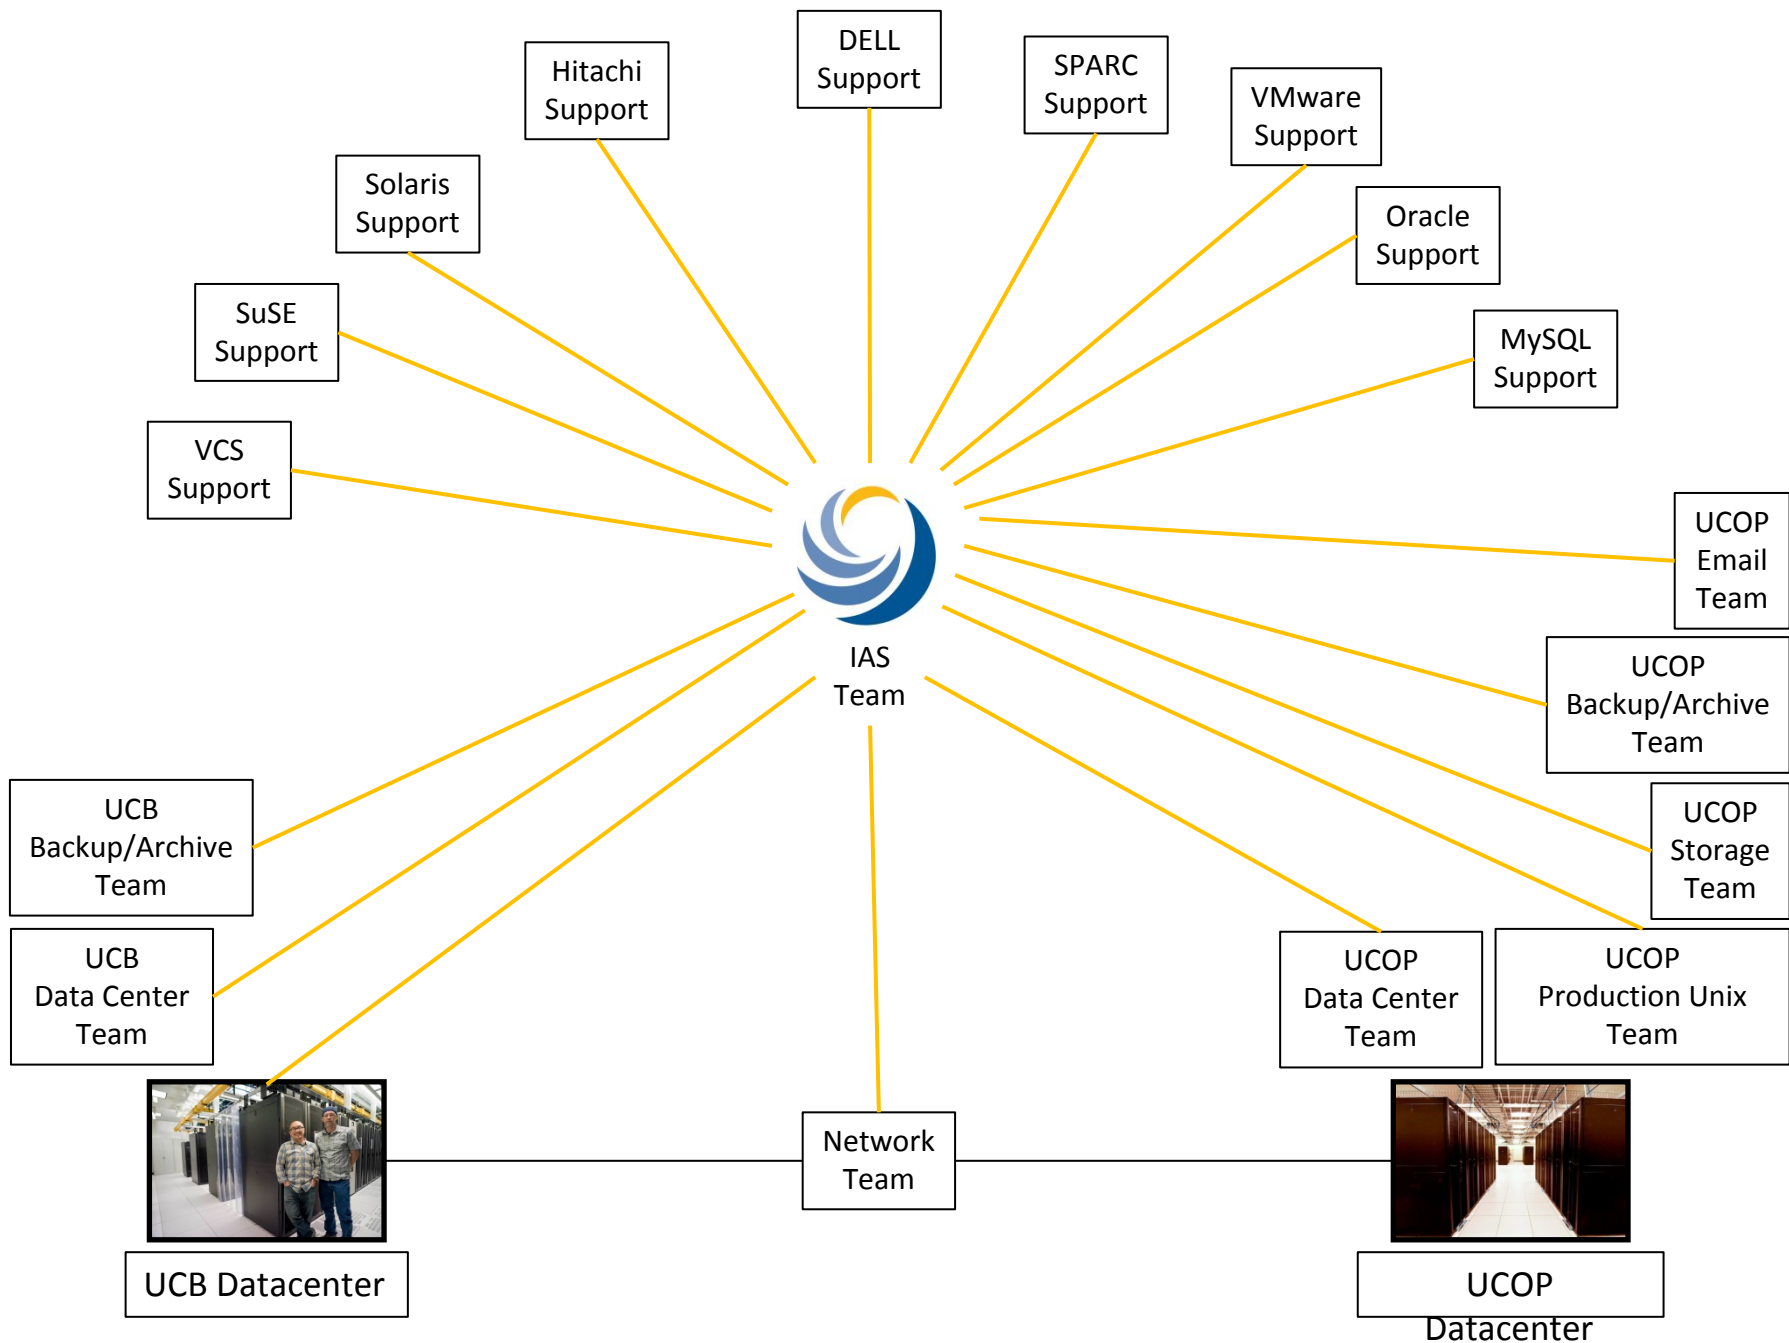


“Overall, I am confident that we can support our entire technology infrastructure in AWS for a cost that is *no greater* than what we currently pay in annual costs for our existing infrastructure.”

- *Strategic Plan for CDL Technology Infrastructure, 2014-2017, dated 6/10/2014*









S3 (Storage)



Route53 (DNS)



EC2 (Servers)



VPC (Networking)

1 Vendor



CloudWatch (Logging)



RDS (Databases)

1 Account



Cloudtrail (Auditing)



ELB (load balancing)



AWS Linux (OS)

What's Next (from UCCSC 2015)

- ~~Complete the migration~~
- Move from On-demand to Reserved instances
- Move all appropriate services to an H/A model
- Create DR capability in an alternate region
- ~~Evaluate EFS as H/A NFS for Homesrv~~
- ~~Evaluate OpsWorks to replace Puppet~~
- ~~Implement 3rd party tools for user access to CloudWatch performance data~~
- ~~CloudTrail notifications from CloudWatch~~
- ~~Look at new databases (Aurora, Redshift, DynamoDB)~~
- Elastic Beanstalk pilot underway (IaaS > PaaS)

What's Next (from UCCSC 2015)

- ~~Complete the migration~~
- **Move from On-demand to Reserved instances**
- Move all appropriate services to an H/A model
- Create DR capability in an alternate region
- ~~Evaluate EFS as H/A NFS for Homesrv~~
- ~~Evaluate OpsWorks to replace Puppet~~
- ~~Implement 3rd party tools for user access to CloudWatch performance data~~
- ~~CloudTrail notifications from CloudWatch~~
- ~~Look at new databases (Aurora, Redshift, DynamoDB)~~
- Elastic Beanstalk pilot underway (IaaS > PaaS)



180 instances
3 availability
zones

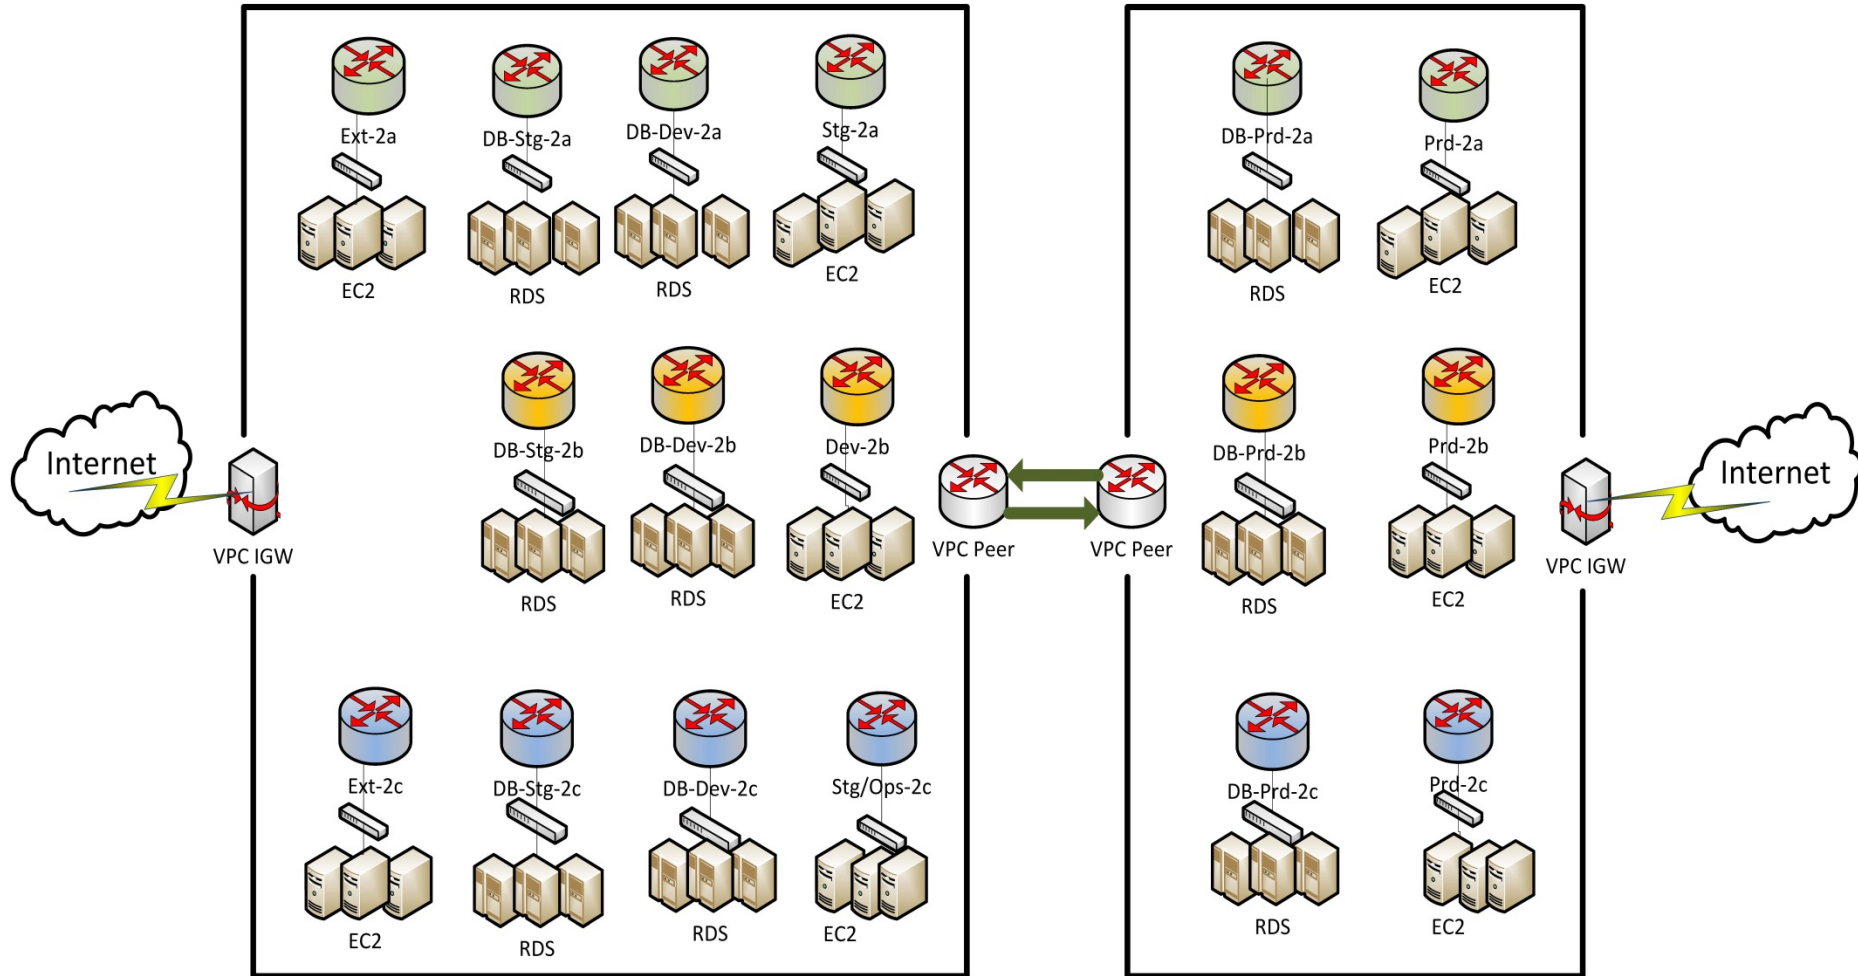
60 instances per
AZ
10 instance types
6 instances per
type

4 reserved instances
2 on-demand instances
6 instances (and a small cost
savings)

AWS Oregon Region

Dev/Stg VPC

Prd VPC



Note: Suffix of 2a(green), 2b(yellow) and 2c(blue) indicates Availability Zones.

What's Next (from UCCSC 2015)

- ~~Complete the migration~~
- Move from On-demand to Reserved instances
- **Move all appropriate services to an H/A model**
- Create DR capability in an alternate region
- ~~Evaluate EFS as H/A NFS for Homesrv~~
- ~~Evaluate OpsWorks to replace Puppet~~
- ~~Implement 3rd party tools for user access to CloudWatch performance data~~
- ~~CloudTrail notifications from CloudWatch~~
- ~~Look at new databases (Aurora, Redshift, DynamoDB)~~
- Elastic Beanstalk pilot underway (IaaS > PaaS)

What's Next (from UCCSC 2015)

- ~~Complete the migration~~
- Move from On-demand to Reserved instances
- Move all appropriate services to an H/A model
- **Create DR capability in an alternate region**
- ~~Evaluate EFS as H/A NFS for Homesrv~~
- ~~Evaluate OpsWorks to replace Puppet~~
- ~~Implement 3rd party tools for user access to CloudWatch performance data~~
- ~~CloudTrail notifications from CloudWatch~~
- ~~Look at new databases (Aurora, Redshift, DynamoDB)~~
- Elastic Beanstalk pilot underway (IaaS > PaaS)



6/1/2016

CDL Infrastructure Migration - 1 Year Later

What's Next (from UCCSC 2015)

- ~~Complete the migration~~
- Move from On-demand to Reserved instances
- Move all appropriate services to an H/A model
- Create DR capability in an alternate region
- ~~Evaluate EFS as H/A NFS for Homesrv~~
- ~~Evaluate OpsWorks to replace Puppet~~
- ~~Implement 3rd party tools for user access to CloudWatch performance data~~
- ~~CloudTrail notifications from CloudWatch~~
- ~~Look at new databases (Aurora, Redshift, DynamoDB)~~
- **Elastic Beanstalk pilot underway (IaaS > PaaS)**

Next Steps for FY1617

- Complete Phase 1 of Disaster Recovery project
- More delegation of responsibilities (instance stop/start for “On-Demand” infrastructure)
- Security self-assessment
- Learn Beanstalk and the Container Service



MIRBANA



Relentless
Innovation



The Fun Never Stops



Changes

NEXT EXIT



Thanks!



Questions?