CDL AWS Environment Orientation

AWS Resources

> *AWS Service Health Dashboard http://status.aws.amazon.com/* (Click the RSS icon for history)
>
> *Amazon Linux AMI Security Center: https://alas.aws.amazon.com/*

General Guidelines

> *In general, the CDL AWS infrastructure consists of virtual machines referred to as EC2 instances. An instance is requested using the online web form on IntraCDL [EC2, RDS]. The CDL AWS environment is created in the AWS us-west-2 region (Oregon), which consists of 3 interconnected availability zones (us-west-2a, us-west-2b, us-west-2c). Spreading services across multiple availability zones provides high availability by reducing the impact of a failure of any single availability zone or component. The AWS environment can be accessed from trusted UCOP networks; see Network User Guide and/or the Internet via the Bastion Service; see AWS Bastion Service Users Guide. User accounts are managed by LDAP (see LDAP Users Guide) which is separate from UCOP AD. Users in AWS have a single NFS home directory that is shared between Development, Stage and Production environments. All AWS infrastructure resources are created and managed in partnership with programs by the IAS (Infrastructure Application/Support) team.*

EC2 Instances

> *The CDL default instance type is a t3.small, with 1 vCPU and 2GB of memory. Other instance types can be requested to meet a particular technical requirement, but will require review and approval by IAS. The target for delivering an instance request is 5 days. The default instance will be delivered running AWS Linux (which is a CentOS variant) and will have a standard set of packages that comes with the AMI (Amazon Machine Image).*
>
> *Additional packages can be installed by the instance owner from the AWS repository using "sudo yum" and if the desired package is not there, can make a request via cdlsys ticket (cdlsys@cdlib.org). IAS staff will open a request to AWS support to add the package to the repository. Instance owners can always download and build their own packages from open source, and install it within their application space on an AWS instance.*
>
> *Instances are named according to the following convention: <program name>-<service name>-<environment name>. For example, ias-nagios2-ops or d2d-htmm-dev. Users will access their instances using SSH. Auto-home has been configured so that a home directory will be automatically mounted when a user logs in to an AWS instance. Permissions to AWS resources like S3 Buckets, CloudSearch, etc. are managed by IAM (Identity and Access Management) roles and policies applied to the EC2 instance. EC2 instances can be stopped/started and rebooted; see Roles and Responsibilities section below.*

EBS Storage

Each instance will be configured with SSD local storage. The original standard magnetic HDD volume type has been deprecated. The minimum storage size for st1 throughput optimized volume type is 500gb. Other storage types can be requested to meet a particular technical requirement, but will require review, approval by IAS and may affect delivery time. Service-specific applications will be installed under the /apps/<role-account> directory.

Encryption: *Werner Voegels, the AWS CTO, has worn a T-shirt on stage that articulates AWS's position on the subject of encryption, which reads: "Encrypt Everything". It's good advice, but when it comes to encryption at rest of EBS volumes, we've been following the advice from a similar T-shirt, worn on stage by Corey Quinn of the "Last Week in AWS" newsletter, which reads: "Encrypt Everything (unless it's hard)". Fortunately, we've come up with a way where it's no longer all that hard. Here's our position on the matter:*

1. *As of 10/11/2021, all new EC2 instances are created with all volumes (system and application) encrypted at rest.*
2. *Pre-existing non-root EBS volume(s) (i.e. "application" volumes) will be migrated over time to encrypted storage in the background, with no downtime for most hosts. We intend to complete this by the end of December 2021.* Exception*: large volumes (>= 1 TB) will require a short downtime, as the background process takes excessively long for these.*
3. *Pre-existing root EBS volumes (i.e. "system" volumes) will remain unencrypted except by request, since encrypting these requires system downtime. We anticipate a migration of our Amazon Linux 2 instances (which, at the time of this writing, is essentially all of them) to Amazon Linux 3 in the not too distant future, and we will be building new Amazon Linux 3 hosts with encrypted EBS volumes to replace these Amazon Linux 2 instances as part of that process.*
4. *All snapshots and DR AMI backups of encrypted volumes will also be encrypted.*

Networking

- *The CDL environment is built using Virtual Private Clouds (VPCs). There is one VPC for Development/Stage and one for Production.*
- *Each instance will have a private IP address (available from other AWS instances or services). A public-facing instance (application server) will be assigned an Elastic IP (EIP, which is a static IP address that can be reached via the Internet).*
- *Both VPCs share a common subnet for EIPs, and have several subnets for internal IP addresses.*
- *Every EC2 instance has a firewall known as a Security Group.*
  - *By default, AWS instances are accessible via ssh from trusted UCOP networks (CDL offices and VPN).*
  - *Internal AWS network traffic is open to all.*
- *Standard port forwarding will be enabled for http and https. Port 80 is port forwarded to 18880 and port 443 is port forwarded to 18443.*
- *DNS is managed in Route53.*

- *For additional details, see the Network User Guide.*

Load Balancing

Application Load Balancers (ALBs) can be used to provide high-availability and variable resources for CDL services that require this capability. ALBs require a more complicated security certificate configuration, and the IP addressing is handled differently, as well. ALBs are available upon request, and will require further consultation with IAS before implementation.

S3 Storage

Two S3 buckets have been created for shared use within a program, one with versioning turned off (the default) and one with versioning turned on. These are named using the following convention: <program>-s3-<environment> and <program>-s3ver-<environment>. Each bucket will have a folder for each service, and within that a subfolder called “glacier” that has a policy which will put the data directly into lower-cost Glacier storage for archival purposes. Additional buckets can be created upon request, after review and approval by IAS.

For additional information, see the S3 and Glacier Storage User's Guide.

EFS (Elastic File System)

EFS (Elastic File System) is available for applications that need dynamic storage (provisioned only as needed) or shared storage across multiple EC2 instances. Note that EFS performance is not comparable to EBS storage and may not be suitable for all workloads.

- *https://aws.amazon.com/efs/features/*

EFS is designed to be highly available and durable. ⚠️ *EFS file systems are not backed up.*

Monitoring and Performance

> *Application and Infrastructure monitoring/performance is provided by the following tools:*
>
> - *Nagios (Nagios User's Guide)*
> - *Librato (Librato / AWS Metrics User's Guide)*
> - *Custom Application Monitoring (Custom Metrics User Guide)*

FTPS and SFTP

> *IAS provides ftps and sftp services for CDL programs and their partners.*
>
> - *FTPS (FTPS Users Guide)*
> - *SFTP (SFTP Users Guide)*

SSL Certificate Management

> *IAS provides SSL certificate management for CDL programs. See SSL Certificate Management Users Guide.*

Accounts & Privileges

> *UNIX user accounts within the AWS environment are maintained in LDAP and not in UCOP AD. Any work below will need to have a cdlsys ticket opened. See your tech lead for more information.*
>
> - *To create a new LDAP account, you will need to work with your tech lead to create access and have it approved.*
> - *Reset your password if you've forgotten or you want it changed.*
> - *sudo privileges and group ownership changes; this will need a tech lead approval.*
>   - *Role accounts are assigned elevated privileges via sudo within the operating system.*

Backup & Recovery

> *Backup and recovery is provided by daily snapshots of **all instances** and is pre-configured during the instance deployment process by IAS. These snapshots are created at 12:01am every night, retained for 35 days, and then deleted. Snapshots on the 1st of the month are retained for 6 months. Additional manual snapshots are available upon request, and can be maintained for longer periods.*
>
> *Images of **production instances** are copied to an alternate AWS region once a week as a recovery point in the event an entire AWS region is unavailable for an extended period of time. IAS maintains backups of all AWS configuration information. Many resources, including IAM, security groups, S3 bucket configurations, Route 53 DNS records, and CloudFormation templates are under version control in a local git repository (itself backed up in an alternate region).*

Databases

> *By default, CDL services will use the managed database service from Amazon called RDS (Relational Database Service). An IAS DBA will create and maintain the RDS instances, while AWS will provide backups, replication, and upgrades. The IAS DBA will work with each service to create new databases or migrate existing ones. See RDS User Guide.*

Email

> *AWS hosts are not configured to receive mail. Outgoing mail can be sent using Amazon Simple Email Service (SES) depending on the application requirements. See Email User Guide for more information.*

Roles & Responsibilities

Application developers are given Unix credentials to access the system, and role accounts that will provide elevated privileges, while the IAS team has AWS credentials to manage the AWS services. An immediate difference is that the service owner is granted the rights and responsibility to patch (and reboot if necessary) their instance; the ability to stop/start the instance via an admin host owned by the program. Service owners are expected to monitor security announcements and patches for third party application software installed locally and patch in a timely manner. Some patches (WordPress, Drupal, PHP) may need to be applied within hours of release.

- *See the Yum User's Guide for more information and policies for system patching.*
- *See the Manual EC2 Stop, Start User's Guide for more information on stop/start of your EC2 instance.*
- *Some AWS privileges are delegated to an aws-ops instance for each program (for example, pub-aws-ops). See AWS CLI User Guide for more information.*

External Accounts

Use Cases

Separate AWS accounts may be created for a program or service. Candidates are AWS services with ephemeral resources ("cattle" not "pets"), permissions that cannot be easily managed with IAM policies, or services that require AWS console access. For example:

- *ElasticBeanstalk*
- *Lambda*

Other services may be enabled based on the application architecture, for example:

- *CloudFront*

External accounts are managed from a master account using AWS Organizations with specific services enabled. External accounts are generally isolated environments. IAS staff will have access to your AWS account. You will have access to billing information for cost awareness. Over time, we will be establishing additional requirements for separate accounts as necessary to ensure the safety and security of the CDL cloud infrastructure as a whole.

IAS Responsibilities and Services Provided

- *Account setup, including AWS console access for designated developers. AWS Console access for interdependent accounts is limited to CDL developers (not external collaborators).*
- *Creation of a VPC peered with the main account, with bi-directional DNS resolution. This allows access to specific resources in the main account, including designated RDS databases, as well as access from the main account to the external account over private IP addresses.*
- *Standard security groups (CDL AWS Standard allowing SSH from trusted sources including all hosts in the main account, HTTP(S) from ANY, HTTP(S) from UCOP).*
- *An EC2 keypair (stored on the program's aws-ops instance), allowing SSH access to EC2 instances in the external account.*
- *Limited log review/monitoring (CloudTrail, Config, Guard Duty and Firewall Manager).*
- *Nagios service monitoring (Nagios host-based monitors, including disk space and load, are not supported).*
- *Route 53 DNS in the main account.*
- *SES verified domains and senders are configured in the main account.*
- *Best effort supporting/troubleshooting issues with individual Amazon Linux hosts or AWS services.*

Services not Provided in External Accounts

- *Backup/Disaster Recovery*
- *Librato performance monitoring (AWS CloudWatch is available)*
- *Unix log review*
- *Syschanges (file integrity monitoring)*
- *Puppet*
- *Unix account management*
- *Access to NFS home directories in the main account*
- *Patching*
- *EIPs (static IP addresses for long-lived EC2 instances)*
- *Databases are provisioned in the main account, with access from external accounts as needed.*

Usernames and Passwords

- *AWS IAM usernames are the same as LDAP (Unix) usernames*
- *Password requirements: Minimum 8 characters, at least one each upper, lower, digit, special character (! @ # $ % ^ & * ( ) _ + - = [ ] { } | ')*
- *All IAM user accounts must have MFA enabled.*

Expectations for Developers

- *If an EC2 instance is running longer than 7 days, it must be patched in accordance with the CDL patching policy.*
- *Any access keys must be rotated on a regular schedule (IAS will calendar the key rotation and remind you).*
- *No AWS credentials are to be placed in an external repository (public or private).*
- *Resources must be tagged for cost reporting across accounts: Program, Service, Environment. (See AWS CLI User Guide and consult with IAS on standards.)*
- *S3 buckets may not be open for public access (read or write).*

Working Across Accounts

- *In general, access from an external account to the main account is limited to resources like RDS databases which are managed in the main account.*
- *SSH access will generally not be allowed from external accounts to the main account; instead SSH from the main account to the external account. You can manage the security groups on resources in the external account to allow any required access from the main account.*
- *S3 bucket access across accounts will be reviewed on a case-by-case basis.*

End Document